arXiv:1410.0986v2 [math.NA] 14 Oct 2014

FAST ALGORITHMS FOR HYPERSPECTRAL DIFFUSE OPTICAL TOMOGRAPHY

ARVIND K. SAIBABA *, MISHA KILMER†, ERIC L. MILLER*, AND SERGIO FANTINI‡

Abstract.
The image reconstruction of chromophore concentrations using Diffuse Optical Tomography (DOT) data can be described mathematically as an ill-posed inverse problem. Recent work has shown that the use of hyperspectral DOT data, as opposed to data sets comprising of a single or, at most, a dozen wavelengths, has the potential for improving the quality of the reconstructions. The use of hyperspectral diffuse optical data in the formulation and solution of the inverse problem poses a significant computational burden. The forward operator is, in actuality, nonlinear. However, under certain assumptions, a linear approximation, called the Born approximation, provides a suitable surrogate for the forward operator, and we assume this to be true in the present work. Computation of the Born matrix requires the solution of thousands of large scale discrete PDEs and the reconstruction problem, requires matrix-vector products with the (dense) Born matrix. In this paper, we address both of these difficulties, thus making the Born approach a computational viable approach for hyperspectral DOT (hyDOT) reconstruction. In this paper, we assume that the images we wish to reconstruct are anomalies of unknown shape and constant value, described using a parametric level set approach, (PaLS) [1] on a constant background. Specifically, to address the issue of the PDE solves, we develop a novel recycling-based Krylov subspace approach that leverages certain system similarities across wavelengths. To address expense of using the Born operator in the inversion, we present a fast algorithm for compressing the Born operator that locally compresses across wavelengths for a given source-detector set and then recursively combines the low-rank factors to provide a global low-rank approximation. This low-rank approximation can be used implicitly to speed up the recovery of the shape parameters and the chromophore concentrations. We provide a detailed analysis of the accuracy and computational costs of the resulting algorithms and demonstrate the validity of our approach by detailed numerical experiments on a realistic geometry.

1. Introduction. Diffuse optical tomography (DOT) is an imaging technique that uses near infrared light to image highly scattering media. A good review has been provided in [3] and an updated version is provided in [2]. The imaging modality has shown great promise as a low-cost alternative or complement to existing medical imaging technology particularly in brain imaging and breast cancer detection. The region of interest is illuminated with near infrared light over a collection of wavelengths and the data are comprised of observations of the resulting scattered diffuse fields at a number of locations surrounding the medium. Given these measurements as well as the partial differential equation governing the interaction of light and tissue (typically, the diffusion equation), we seek to recover space and time-varying maps (i.e. images) of concentrations of physiologically relevant chromophores such as oxygenated and deoxygenated hemoglobin (HbO_2 and HbR respectively), lipid, and water (H_2O) as well as properties governing the scattering of light within the medium.

The recovery of images of chromophore concentrations can be mathematically posed as a nonlinear inverse problem. However, due to the diffusive physics associated with this problem as well as limitations concerning the geometric distribution of sources and detectors, image recovery is an ill-posed inverse problem. New technology developed in our research group allows for the collection of hyperspectral data (over 100 bands in the near infrared portion of the spectrum). Although we have demonstrated [25, 26, 27] that the availability of more information using multiple wavelengths increases the accuracy of the reconstruction, the use of hyperspectral data poses a significant computational burden in the context of image recovery. We are interested in developing computationally efficient methods for hyperspectral diffuse optical tomography (HyDOT) with specific application towards breast imaging in which the breast is placed in between two parallel plates.

To motivate the need for fast algorithms for hyperspectral DOT, we outline here the expected costs in terms of storage and computation. To make ideas concrete, we consider an experimental setup for detecting tumors in breast tissue (see Figure 1.1). We use N_s near-infrared sources to illuminate the medium of interest. The sources are constrained to lie on the top plane and detectors are constrained to be on a different plane so that for a given source, we have N_{ds} detectors measuring photon fluence at N_λ wavelengths. This results in $M = N_s N_{ds} N_\lambda$ measurements. We also assume that the domain has been discretized into a grid with N vertices.

*Department of Electrical and Computer Engineering, Tufts University {arvind.saibaba,eric.miller}@tufts.edu
†Department of Mathematics, Tufts University misha.kilmer@tufts.edu
‡Department of Biomedical Engineering, Tufts University sergio.fantini@tufts.edu

Number of	Symbol	Typical number
Sources	N_s	$10 - 100$
Detectors / source	N_{ds}	$3 - 10$
Wavelengths	N_λ	$10 - 200$
Grid size	N	$32^3 - 100^3$

Table 1.1: Typical range of parameters for the hyperspectral DOT problem

In Table 1.1, we have listed the range of various parameters that one might encounter in practice. Since the most accurate forward model in terms of the unknown voxel values is nonlinear, standard numerical/optimization approaches to solve the inverse problem repeatedly linearize the forward problem about a current estimate [3, 12]. Each optimization step then requires the solution of the forward and the adjoint PDE for each source-detector set [3]. In all, we need to solve $N_s(N_{ds}+1)N_\lambda$ systems of equations *at each optimization step*, which amounts to about 2×10^5 systems of equations for the range of parameters described in Table 1.1. Thus, for finely discretized fields, even with a solver of optimal complexity $\mathcal{O}(N)$, use of the nonlinear forward model poses a significant computational challenge because the resulting cost is $\mathcal{O}(N_s(N_{ds}+1)N_\lambda N)$ flops.

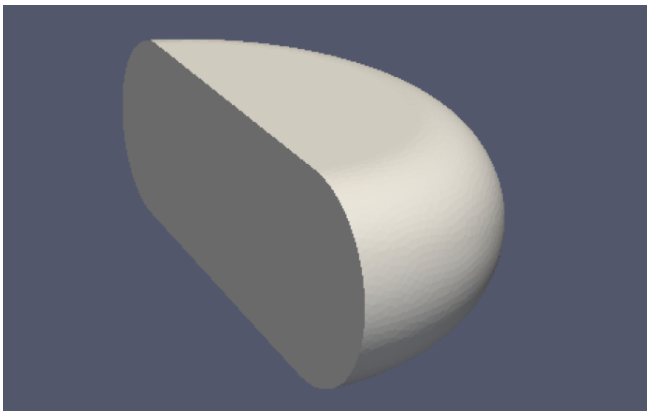

Fig. 1.1: The breast shaped phantom used as the imaging medium. At its widest, the phantom is 12 cm long and 5 cm thick.

To mitigate the computational burden, in the present work we restrict ourselves to the case where the absorption can be represented as a small perturbation about the background medium. Then, we can use the Born approximation to linearize the forward model. Although this approximation introduces a modeling error and has known limitations [7], studies with experimental data have validated the utility of this model for hyperspectral DOT [25, 26, 27]. However, since the domain of interest has a complex shape, the Green's function required for the construction of the linearized operator is not known in closed form. Therefore, we need to use a PDE formulation to compute the incident and adjoint fields. An advantage of the PDE formulation is that we can handle known, non-constant background which may be obtained by imaging using a high resolution modality (such as magnetic resonance imaging, MRI) along with image segmentation [8]. Although the breast is a highly heterogenous medium, based on previous work we model it as a piecewise constant medium with homogenous background. For example, the authors in [36] use prior anatomical knowledge to derive a piecewise constant medium. Although the Born approximation reduces the computational burden associated with the imaging problem tremendously, computing the Born operator $\mathbf{H}$ (which is a linear mapping between the perturbation of interest and the measurements) still requires the solution of many large-scale discretized PDEs for the incident and adjoint field; in sum, about $\mathcal{O}(N_s(N_{ds}+1)N_\lambda)$ large scale, sparse linear systems corresponding to the discrete diffusion equation.

Furthermore, the cost to store $\mathbf{H}$ and form matrix-vector products (matvecs) with $\mathbf{H}$ during the optimization that are needed to invert for the desired parameters scales as $\mathcal{O}(MN)$, where M is the number of measurements and N is the number of unknowns. Therefore, we develop a recursive algorithm to generate a low-rank approximation to $\mathbf{H}$ and use this approximation in the optimization. If the numerical rank of the low-rank representation is R then the cost of storing factors and forming matvecs are $\mathcal{O}(R(M+N))$. Of course, the optimal rank-R approximation could in theory be computed from the SVD of $\mathbf{H}$ if we had $\mathbf{H}$. This is inefficient on two fronts: a) it requires the full computation and storage of as well as multiple accesses to the very large, dense matrix $\mathbf{H}$ b) the cost of an SVD on $\mathbf{H}$, if we had it, is still prohibitively expensive at $\mathcal{O}\left(\min\{M,N\}^2, \max\{M,N\}\right)$ flops. On the other hand, our algorithm, because it operates in a spatially recursive manner requiring, at the finest scale, only local information, and using rank revealing factorizations to aggregate information globally, does not require that $\mathbf{H}$ be fully stored.

We have three main goals in this paper. The first is to design an algorithm to overcome the challenge of computing the solution of $\mathcal{O}(10^5)$ number of large-scale parametric linear systems needed to obtain certain information necessary to compute our approximate Born matrix. The second is to design an efficient algorithm to compute and store a meaningful low-rank representation of the measurement operator $\mathbf{H}$. The third is to use this low-rank representation for recovering the parameters that define our chromophore image. We now summarize the key components of our fast algorithm and contributions in this paper:

- For each source-detector set, we need to compute the incident field and the adjoint field for hundreds of wavelengths. We derive a novel Krylov recycling subspace approach to solve the corresponding systems of large-scale parametric linear systems which takes advantage of similarities in the systems across wavelengths. This is described in detail in Section 3.
- For the problem at hand, the sources, detectors and the perturbation are well-separated from each other. Under these circumstances there is strong numerical evidence that the Born operator $\mathbf{H}$ is low-rank, a feature that has been exploited to achieve computational savings [9]. In this work, we give a new approach for computing a low rank approximation, and use this in our numerical results on the parametric inverse problem. The storage of the matrix $\mathbf{H}$ in its entirety is avoided; instead, the rows of $\mathbf{H}$ corresponding to a single source-detectors set are constructed, compressed and then recursively compressed across multiple source-detectors pairs to obtain a low-rank factorization $\mathbf{H} \approx \hat{\mathbf{H}} = \mathbf{U}\mathbf{V}^T$. We provide a detailed outline of the algorithm and analysis of the computational costs in Section 4. Because we fully compute the entries of $\mathbf{H}$, we need the fast Krylov solvers because we need to solve $\mathcal{O}(10^5)$ systems.
- Based on recent success in the context of DOT, we employ the parametric level set (PaLS) approach to represent the chromophore image we want to recover. The PaLS approach, developed for inverse problems in [1] and subsequently applied to diffuse optical tomography [27, 25], has the advantage of explicitly describing the geometry of the anomaly.
  As a result, the recovery of the chromophore image is obtained by solving a nonlinear least-squares solver on a problem in which we have replaced the Born operator by the approximation described above. We show how to accelerate the reconstruction of the chromophore concentration and shape parameters by using the compressed measurement operator $\hat{\mathbf{H}}$ is used in place of the full measurement operator. Error bounds are provided based on the error in the low-rank representation and are useful for the optimization routine. This is outlined in Section 5.

Related work: We briefly review other approaches to solve the parametric system of equations that represent the propagation of light in diffuse medium. One approach to deal with parametric system of equations is the use of spectral methods by expanding the matrix coefficients in a series of orthogonal polynomials (such as Chebyshev, Legendre, etc.) and solving for the coefficients of the orthogonal polynomials using a Collocation method or a residual minimizing Galerkin projection method (see [11] and references therein). The low-rank property of the solutions arising from system of parametric coefficients with coefficients smoothly dependent on the parameters was demonstrated by Kressner and Tobler [24] and they developed global Krylov subspace algorithms that exploited the low-rank nature to minimize computational and storage costs. In both approaches, one has to solve a coupled system of equations. This can be computationally expensive and, in order to ensure rapid

convergence, a preconditioner that is effective across all the shifts is necessary but difficult to choose in practice. Other approaches to deal with the expensive cost of solving parametric linear systems is to use parametric model reduction which is reviewed in [6].

Other works have also considered the compression of the measurement operator $\mathbf{H}$ that maps the perturbation to the measurements. In [31], the authors develop analytical formulas for inversion based on Fourier analysis when the sources and detectors are distributed uniformly on the boundary of a regular geometry such as a plane, cylinder, or sphere. In our previous work, we have exploited the structure of the Green's function in regular geometries to decompose the Born operator into a number of sparse easily computed matrices [21]. The approach of compressing the operator $\mathbf{H}$ is similar to that derived in [9]. Here, the authors consider compression across multiple sources and detectors for a given frequency using randomized SVD and then recursively compressing the low-rank factors across multiple frequencies. However, we cannot adopt their framework directly for the following reasons. Firstly, in our system, the detectors are not shared across all of the sources. Consequently, the pre-processing step that compresses the incoming field and the data, as described in [9], cannot be used directly. Secondly, the authors in [9] advocate compression of the measurement operator across multiple sources and detectors for a given frequency and then recursively combining the low-rank factorizations across different wavelengths. Memory limitations prevent computations of the entire measurement operator in its entirety. Because of the way our computations are organized, we choose to first compress across multiple wavelengths and detectors for a given source and then combine the low rank factorizations across multiple sources.

2. Forward Problem. In this section, we give the specifics of the forward problem and discretization used for the associated PDEs.

2.1. Governing equations. The radiative transport physics associated with the propagation of light through a medium can be approximated by the diffusion model of the form in the domain Ω [3]

$$-\nabla \cdot D^{\text{tot}}(\mathbf{r},\lambda)\nabla\phi(\mathbf{r},\lambda) + \nu\mu_a^{\text{tot}}(\mathbf{r},\lambda)\phi(\mathbf{r},\lambda) = S(\mathbf{r},\lambda) \qquad \mathbf{r}\in\Omega \tag{2.1}$$

$$\phi(\mathbf{r},\lambda) = 0 \qquad \mathbf{r}\in\partial\Omega_D \tag{2.2}$$

$$\phi(\mathbf{r},\lambda) + 2AD^{\text{tot}}(\mathbf{r},\lambda)\frac{\partial\phi(\mathbf{r},\lambda)}{\partial n} = 0 \qquad \mathbf{r}\in\partial\Omega_R \tag{2.3}$$

where $D^{\text{tot}}(\mathbf{r},\lambda)$ is the diffusion coefficient and is related to the reduced scattering coefficient $\mu_s'(\lambda)$ as $D(\mathbf{r},\lambda) = \nu/3\mu_s'(\mathbf{r},\lambda)$. We also denote by $\partial\Omega_D$ the portion of the boundary over which zero Dirichlet boundary conditions are imposed (curved boundaries and chest wall) and $\partial\Omega_R$ corresponds to the boundary portion over which Robin boundary conditions are imposed corresponding to a refractive index mismatch (top and bottom flat regions). The coefficient A is a function of the refractive index of the medium. We denote by $\phi(\mathbf{r},\lambda)$ the photon fluence at a position $\mathbf{r}$ due to a source of wavelength λ injected into the medium, and ν is the electromagnetic propagation velocity within the medium. Further, $\mu_a^{\text{tot}}(\mathbf{r},\lambda)$ is the absorption coefficient. The quantity $S(\mathbf{r},\lambda)$ is the photon source with units of optical energy per unit time per unit volume and typically written in terms of a delta function; that is, $S(\mathbf{r},\lambda) = S_0(\lambda)\delta(\mathbf{r}-\mathbf{r}_s)$, with $S_0(\lambda)$ the source power at wavelength λ.

We decompose the absorption, $\mu_a^{\text{tot}}(\mathbf{r},\lambda)$, into a constant background absorption $\mu_a(\lambda)$ and a spatially varying perturbation $\Delta\mu_a(\mathbf{r},\lambda)$. The total fluence, ϕ, is decomposed into an incident field ϕ_i and a scattered field ϕ_s. Likewise, we can expand the diffusion $D^{\text{tot}}(\mathbf{r},\lambda)$ into the sum of a homogenous background term $D(\lambda)$ and a perturbation $\Delta D(\mathbf{r},\lambda)$. However, the spatial dependence of diffusion is minimal in healthy breasts [37] and spatial contrast in breast tumors is either non-existent or small [16, 15]. To simplify matters we assume that $D^{tot}(\mathbf{r},\lambda)$ is independent of $\mathbf{r}$ and $\Delta D(\mathbf{r},\lambda) = 0$. Therefore $D^{\text{tot}}(\mathbf{r},\lambda)$ can be represented entirely as $D(\lambda)$ and we can then divide throughout by $D(\lambda)$ (see for example [26]). The equation for the incident field ϕ_i and the scattered field ϕ_s can be obtained by substituting $\phi(\mathbf{r},\lambda) = \phi_i(\mathbf{r},\lambda) + \phi_s(\mathbf{r},\lambda)$ and collecting the appropriate terms, and is therefore,

$$-\nabla^2\phi_i(\mathbf{r},\lambda) + \frac{\nu\mu_a(\lambda)}{D(\lambda)}\phi_i(\mathbf{r},\lambda) = \frac{S_0(\lambda)}{D(\lambda)}\delta(\mathbf{r}-\mathbf{r}_s) \qquad \mathbf{r}\in\Omega \tag{2.4}$$

$$-\nabla^2\phi_s(\mathbf{r},\lambda) + \frac{\nu\mu_a(\lambda)}{D(\lambda)}\phi_s(\mathbf{r},\lambda) = -\frac{\nu\Delta\mu_a(r,\lambda)}{D(\lambda)}\left(\phi_i(\mathbf{r},\lambda)+\phi_s(\mathbf{r},\lambda)\right) \qquad \mathbf{r}\in\Omega \tag{2.5}$$

Under the Born approximation, the scattered field is assumed to be much smaller than the incident field, i.e. $\phi_s \ll \phi_i$ and therefore, the total fluence $\phi(\mathbf{r},\lambda) = \phi_i(\mathbf{r},\lambda) + \phi_s(\mathbf{r},\lambda)$ in the right hand side of equation (2.5) can be replaced by $\phi_i(\mathbf{r},\lambda)$. As a result of this approximation, there is a linear relation between the scattered fluence rate $\phi_s(\mathbf{r},\lambda)$ and the perturbation of absorption $\Delta\mu_a(\mathbf{r},\lambda)$.

It should also be noted that both the scattered field $\phi_s(\mathbf{r}\,\lambda)$ and $\phi_i(\mathbf{r},\lambda)$ satisfy the same boundary conditions in the equations (2.2)-(2.3). Furthermore, if additional information such as spatial variability is known about the background properties of diffusion and absorption (currently assumed to be homogenous) they can be incorporated into this model [8]. The solution to the photon fluence ϕ_s computed at the measurement location $\mathbf{r}_d$ for a particular wavelength λ can be written using the following integral equation

$$\phi(\mathbf{r}_d,\lambda) = \phi_i(\mathbf{r}_d,\lambda) + \phi_s(\mathbf{r}_d,\lambda) \quad \approx \quad \phi_i(\mathbf{r}_d,\lambda) - \int_\Omega \phi_d(\mathbf{r},\lambda)\nu\Delta\mu_a(\mathbf{r},\lambda)\phi_i(\mathbf{r},\lambda)d\mathbf{r} \tag{2.6}$$

where $\phi_d(\mathbf{r},\lambda)$, which we call the adjoint field, can be derived using the reciprocity property of the Greens function and satisfies the system of equations along with the same boundary conditions in equations (2.2)-(2.3)

$$-\nabla^2\phi_d(\mathbf{r},\lambda) + \frac{\nu\mu_a(\lambda)}{D(\lambda)}\phi_d(\mathbf{r},\lambda) \quad = \quad \frac{1}{D(\lambda)}\delta(\mathbf{r}-\mathbf{r}_d) \qquad \mathbf{r}\in\Omega \tag{2.7}$$

and $\mathbf{r}_d$ corresponds to the detector location. To relate the scattered fluence to the concentrations of chromophores, the perturbation $\Delta\mu_a(\mathbf{r},\lambda)$ is decomposed in terms of piecewise constant functions as

$$\Delta\mu_a(\mathbf{r},\lambda) \overset{\text{def}}{=} \sum_{l=1}^{N_{sp}} \varepsilon_l(\lambda)c_l\chi(\mathbf{r}) \qquad \chi(\mathbf{r}) = \begin{cases} 1 & \mathbf{r}\in\mathcal{D} \\ 0 & \mathbf{r}\in\Omega\backslash\mathcal{D} \end{cases} \tag{2.8}$$

where N_{sp} is the number of species, ε_l is the extinction coefficient of species l at wavelength λ, c_l is the concentration of species l and χ is an indicator function which depends on $\mathcal{D}$, the domain of support for the perturbation we wish to image. For the purpose of this paper, we will consider that the chromophore concentrations are co-located. This choice was also considered in [27]. Further details regarding the governing partial differential equations can be obtained from the following references [25, 26, 27].

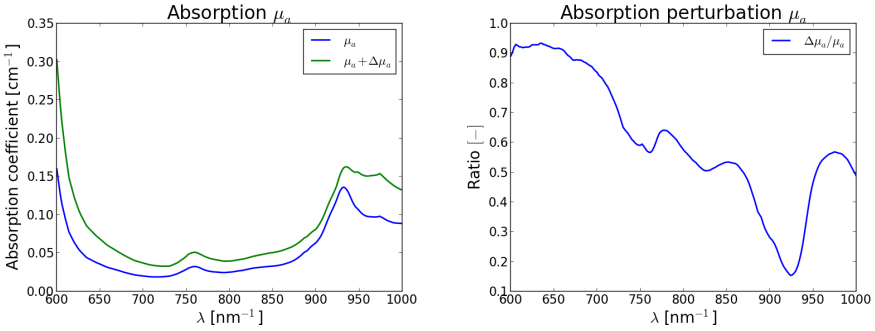

Fig. 2.1: The absorption coefficient μ_a as a function of λ in the range $\lambda \in [600, 1000]$ [nm]. The chromophore concentrations of the background and the perturbation used to generate this plot are provided in Table 6.1.

We also model the diffusion coefficient $D(\lambda)$ using Mie scattering theory [26] as

$$D(\lambda) = \frac{\nu\Psi}{3}\left(\frac{\lambda}{\lambda_0}\right)^b \tag{2.9}$$

The reference wavelength λ_0 is chosen as 600 nm and Ψ has units of cm^{-1}. The scattering pre-factor Ψ depends primarily on the number and size of scatterers, and a scattering exponent b depends on the size of scatterers in the medium [17].

2.2. Discretization using finite elements. To solve the systems of equations (2.5) and (2.4) we use the standard linear Galerkin finite element approach. Expanding the solution field in an appropriately chosen finite dimensional basis $u_m(\mathbf{r})$, i.e.,

$$\phi_i(\mathbf{r},\lambda) \approx \sum_{m=1}^{N} \hat{\phi}_{i,m}(\lambda) u_m(\mathbf{r}) \quad \text{and} \quad \phi_d(\mathbf{r},\lambda) \approx \sum_{m=1}^{N} \hat{\phi}_{d,m}(\lambda) u_m(\mathbf{r})$$

Let us denote the discretized incident photon fluence field as $\boldsymbol{\phi}_i(\lambda) = [\hat{\phi}_{i,1}(\lambda), \ldots, \hat{\phi}_{i,N}(\lambda)]^T$ and the scattered photon fluence $\boldsymbol{\phi}_d(\lambda) = [\hat{\phi}_{d,1}(\lambda), \ldots, \hat{\phi}_{d,N}(\lambda)]^T$. The resulting system of equations can be summarized as

$$\left(\mathbf{K} + \frac{\nu\mu_a}{D}(\lambda)\mathbf{M} + \frac{1}{2AD(\lambda)}\mathbf{R}\right)\boldsymbol{\phi}_i(\lambda) = \frac{1}{D(\lambda)}\mathbf{b}_i \tag{2.10}$$

where the matrices $\mathbf{K}$, $\mathbf{M}$ and $\mathbf{R}$ have entries given by

$$\begin{aligned} \mathbf{K}_{jk} &= \int_\Omega \nabla u_k(\mathbf{r}) \cdot \nabla u_j(\mathbf{r}) d\mathbf{r} \qquad \mathbf{R}_{jk} = \int_{\partial\Omega_R} u_j(\mathbf{r}) u_k(\mathbf{r}) d\mathbf{r} \\ \mathbf{M}_{jk} &= \int_\Omega u_k(\mathbf{r}) u_j(\mathbf{r}) d\mathbf{r} \end{aligned} \tag{2.11}$$

for $j,k = 1,\ldots,N$ and the vector $\mathbf{b}_i$ has entries $\mathbf{b}_{i,j} = \int_\Omega \delta(\mathbf{r}-\mathbf{r}_s) u_j(\mathbf{r}) d\mathbf{r}$. A similar equation can be derived for the adjoint field $\boldsymbol{\phi}_d$ with the same matrices and a different right hand side $\mathbf{b}_d$ which has entries $\mathbf{b}_{d,j} = \int_\Omega \delta(\mathbf{r}-\mathbf{r}_d) u_j(\mathbf{r}) d\mathbf{r}$.

The measurements that are collected at the detector are the values of the photon fluence produced as a result of different sources excited at various wavelengths. Let us denote by $\mathbf{y}$ the vector of measurements obtained as

$$\mathbf{y} = \sum_{k=1}^{N_{\text{sp}}} c_i \mathbf{E}_i \mathbf{H} \boldsymbol{\mu} + \boldsymbol{\eta} \qquad \boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \mathbf{W}^{-2}) \tag{2.12}$$

where the rows of $\mathbf{H}$ are constructed by the discretized representation of the integral equation (2.6) and $\boldsymbol{\mu}$ is a discrete representation of the shape of the absorption perturbation $\chi(\mathbf{r})$. The matrices $\mathbf{E}_i = \mathbf{I}_{N_s} \otimes \mathbf{I}_{N_{ds}} \otimes \text{diag}(\varepsilon_k(\lambda))$ for $k = 1,\ldots,N_{\text{sp}}$ and ε_k are the extinction coefficients for species k that is a function of wavelength. Furthermore, c_k are the concentration of the k-th chromophore species. The measurements are typically corrupted by noise which we model as a Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{W}^{-2})$. The noise covariance $\mathbf{W}$ is modeled as a diagonal matrix with diagonal entries $1/\sigma_m$. The relationship between the standard deviation σ_m to the signal to noise ratio (SNR) is described in [26].

3. Fast solvers for the Born approximation. In order to construct the measurement operator $\mathbf{H}$ and thereby solving the inverse problem, we need the solution of the incident field $\boldsymbol{\phi}_i$ and $\boldsymbol{\phi}_d$ corresponding to multiple source and detector locations and multiple wavelengths. In this Section, we will discuss an efficient solver for the computing the sequence of parametrized systems using a novel recycling approach based on Krylov subspaces. Krylov subspace methods are a popular class of algorithms for iterative solution to linear systems. Recycling Krylov subspaces in the context of parametric systems with smoothly varying coefficients has been previously considered in [10, 23, 34].

3.1. Recycling across wavelengths. For simplicity of notation, we denote by

$$\sigma_j \stackrel{\text{def}}{=} \frac{\nu\mu_a(\lambda_j)}{D(\lambda_j)} \qquad \text{and} \qquad \sigma'_j \stackrel{\text{def}}{=} \frac{1}{2AD(\lambda_j)}$$

the shifts corresponding to wavelengths $j = 1,\ldots,N_\lambda$, and by $\mathbf{x}_j \stackrel{\text{def}}{=} \boldsymbol{\phi}_i D(\lambda_j)$ so that equation (2.10)

$$(\mathbf{K} + \sigma_j \mathbf{M} + \sigma'_j \mathbf{R})\mathbf{x}_j = \mathbf{b} \qquad j = 1,\ldots,N_\lambda \tag{3.1}$$

where matrices $\mathbf{K}, \mathbf{M}$ and $\mathbf{R}$ and vector $\mathbf{b}$ are independent of the shift σ and σ'. The solution to the adjoint field $\boldsymbol{\phi}_d$ can be written in a similar fashion with a different right hand side $\mathbf{b}_d$. We first make the following change of variables $\mathbf{K} \leftarrow \mathbf{M}^{-1/2}\mathbf{K}\mathbf{M}^{-1/2}$, $\mathbf{R} \leftarrow \mathbf{M}^{-1/2}\mathbf{R}\mathbf{M}^{-1/2}$ and transform the vectors as $\mathbf{x}_j \leftarrow \mathbf{M}^{-1/2}\mathbf{x}$ and $\mathbf{b} \leftarrow \mathbf{M}^{-1/2}\mathbf{b}$. This can be done efficiently by using a lumped mass matrix [20]. Denoted by $\mathcal{K}_n\{\mathbf{K}, \mathbf{b}\}$, the Krylov subspace of the matrix $\mathbf{K}$ with starting vector $\mathbf{b}$, is defined as

$$\mathcal{K}_n\{\mathbf{K}, \mathbf{b}\} \stackrel{\text{def}}{=} \text{Span}\{\mathbf{b}, \mathbf{K}\mathbf{b}, \ldots, \mathbf{K}^{n-1}\mathbf{b}\}$$

Krylov subspaces enjoy an interesting property called shift invariance [38], i.e.

$$\mathcal{K}_n\{\mathbf{K}, \mathbf{b}\} = \mathcal{K}_n\{\mathbf{K} + \sigma\mathbf{I}, \mathbf{b}\}$$

Several efficient methods exist for solving the system of equations (3.1) (if $\mathbf{R} = 0$ corresponding to Neumann b.c.s), which solve for multiple shifts roughly at the cost of solving a single system. This is accomplished by generating a subspace that is independent of the shift and use the shift-invariant property of Krylov subspaces (for a detailed review, see [38, Section 14.1] and references therein). However, the presence of a third matrix $\mathbf{R}$ destroys the shift-invariant property of the Krylov subspace methods, unless $\mathcal{K}_n\{\mathbf{K}, \mathbf{b}\}$ is an invariant subspace of $\mathbf{R}$. Even though the shift-invariant property does not hold, *we can utilize information from the solution of the shifted system of equations* $(\mathbf{K} + \sigma_j\mathbf{I})\mathbf{x}_j = \mathbf{b}$ *if the perturbation* $\mathbf{R}$ *is not too large in magnitude.* This is the main idea behind the recycling approach to Krylov subspaces that we are proposing.

We construct a shift-invariant basis $\mathcal{K}_n\{\mathbf{K}, \mathbf{b}\}$ by running n steps of the Arnoldi algorithm and we extract the k smallest Harmonic Ritz eigenvalues and eigenvectors to construct $\mathbf{U}$ and $\mathbf{C}$ both in $\mathbb{R}^{N\times k}$ such that

$$\mathbf{K}\mathbf{U} = \mathbf{C} \qquad \mathbf{C}^T\mathbf{C} = \mathbf{I} \tag{3.2}$$

By using the shift-invariant property of Krylov subspaces, we know that $\mathbf{U}$ is also an approximately invariant subspace of $\mathbf{K}+\sigma_j\mathbf{I}$. We now consider the extension of the solution to the parametric system $\mathbf{A}_j \stackrel{\text{def}}{=} \mathbf{K} + \sigma_j\mathbf{I} + \sigma'_j\mathbf{R}$. When $\|\mathbf{R}\|_2$ is small relative to $\|\mathbf{K}\|_2$, this can be considered a perturbation to the system $\mathbf{K} + \sigma_j\mathbf{I}$ for which we have already generated an approximately invariant subspace. To compute a relation of the kind in Equation (3.2) for the matrix $\mathbf{A}_j$, we proceed as

$$\mathbf{A}_j\mathbf{U} = \mathbf{C} + \sigma_j\mathbf{U} + \sigma'_j\mathbf{R}\mathbf{U} \stackrel{\text{def}}{=} \mathbf{C}'_j$$

In the above equation, the most expensive step is computing $\mathbf{R}_\mathbf{U} \stackrel{\text{def}}{=} \mathbf{R}\mathbf{U}$. The matrix $\mathbf{R}$ only has non-zero entries corresponding to boundaries at which there is refraction index mismatch, in our application it is limited to the top and the bottom boundaries. As a consequence, $\mathbf{R}$ is even more sparse compared to $\mathbf{K}$. Moreover, the matrix $\mathbf{R}\mathbf{U}$ can be precomputed since it will be used across each shift. Compute the thin QR decomposition which using MATLAB notation we represent as $[\mathbf{Q}_j, \mathbf{Y}_j] = \text{qr}(\mathbf{A}_j\mathbf{U}, 0)$. The updated updated matrices can now be computed as $\mathbf{U}_j \stackrel{\text{def}}{=} \mathbf{U}\mathbf{Y}_j^{-1}$ and $\mathbf{C}_j = \mathbf{Q}_j$ and satisfy the relation

$$\mathbf{A}_j\mathbf{U}_j = \mathbf{C}_j \qquad \mathbf{C}_j^T\mathbf{C}_j = \mathbf{I}.$$

Here and henceforth, $\mathbf{U}_j$ is not computed explicitly, rather a solve using the upper triangular matrix $\mathbf{Y}_j$ is performed when it is necessary to form products with $\mathbf{U}_j$.

We adopt the same recycling strategy as [34]. However, where we differ from this approach is the way we construct and update the approximate invariant subspace, as we now explain. The strategy in [34] is to first generate an approximate a recycling basis for the matrix $\mathbf{A}_1$ and then update the recycling subspace of perturbed matrices $\mathbf{A}_j = \mathbf{A}_{j-1} + \Delta\mathbf{A}_j$ after the convergence of the augmented Krylov solver for each of the system. In the context of our problem, note that

$$\Delta\mathbf{A}_j \stackrel{\text{def}}{=} (\sigma_j - \sigma_{j-1})\mathbf{I} + (\sigma'_j - \sigma'_{j-1})\mathbf{R}.$$

Algorithm 1 Augmented GMRES for parametric systems

1: Generate a basis for Krylov subspace $\mathcal{K}_n\{\mathbf{K}, \mathbf{b}\}$

$$\mathbf{K}\mathbf{V}_n = \mathbf{V}_{n+1}\bar{\mathbf{T}}_n \qquad \mathbf{y}_j \stackrel{\text{def}}{=} \arg\min \left\| \mathbf{V}_{n+1}^T\mathbf{b} - \left(\bar{\mathbf{T}}_{n+1} + \sigma_j \begin{bmatrix}\mathbf{I}\\ \mathbf{0}\end{bmatrix}\right)\mathbf{y} \right\|$$

2: Solve the eigenvalue problem $\bar{\mathbf{T}}_n^T\bar{\mathbf{T}}_n\mathbf{z} = \theta\mathbf{T}_n^T\mathbf{z}$ and retain k smallest eigenvalues in magnitude. Collect the eigenvectors into a matrix $\mathbf{Z}_k$ and the corresponding eigenvalues as $\mathbf{\Theta}_k$.
3: Compute $\mathbf{C} = \mathbf{V}_{n+1}\bar{\mathbf{T}}_n\mathbf{Z}_k\mathbf{\Theta}_k$ and $\mathbf{U}_k = \mathbf{V}_{n+1}\bar{\mathbf{T}}_n\mathbf{Z}_k$
4: Compute the thin QR $[\mathbf{C}, \mathbf{Y}] = \text{qr}(\mathbf{C}, 0)$ and set $\mathbf{U} = \mathbf{U}\mathbf{Y}^{-1}$ {Shift-invariant deflation subspace}
5: Compute the initial solution $\tilde{\mathbf{x}}_{0,j} = \mathbf{V}_n\mathbf{y}_j$ for $j = 1, \ldots, N_\lambda$
6: **for** $j = 1, \ldots, N_\lambda$ **do**
7: $\quad\mathbf{C}'_j = \mathbf{C} + \sigma_j\mathbf{U} + \sigma'_j\mathbf{R}\mathbf{U}$ {Compute new deflation subspace}
8: $\quad[\mathbf{C}_j, \mathbf{Y}_j] = \text{qr}(\mathbf{C}'_j, 0)$ and set $\mathbf{U}_j = \mathbf{U}\mathbf{Y}_j^{-1}$.
9: $\quad$Compute $\mathbf{r}_j = \mathbf{b} - \mathbf{A}_j\mathbf{x}_{0,j}$ and $\mathbf{x}_{-1,j} \stackrel{\text{def}}{=} \tilde{\mathbf{x}}_{0,j} + \mathbf{U}_j\mathbf{C}_j^T\mathbf{r}_j$ {Initial deflation}
10: $\quad$Generate $\mathbf{V}_{m-k+1}^{(j)}$ and $\bar{\mathbf{T}}_{m-k}^{(j)}$ by applying $m-k$ steps of the Arnoldi algorithm using matrix $(\mathbf{I} - \mathbf{C}_j\mathbf{C}_j^T)\mathbf{A}_j$ applied to the initial deflated vector $(\mathbf{I} - \mathbf{C}_j\mathbf{C}_j^T)\mathbf{r}_j$
11: $\quad$Solve the least-squares system of equations (3.5) with $\beta_j = \|(\mathbf{I} - \mathbf{C}_j\mathbf{C}_j^T)\mathbf{r}_j\|_2$

$$[\mathbf{Q}, \mathbf{Z}] = \text{qr}(\bar{\mathbf{T}}_{m-k}^{(j)}, 0) \quad \mathbf{y}_{2,j} = \mathbf{Z}^{-1}(\mathbf{Q}^T\beta_j\mathbf{e}_1) \quad \mathbf{y}_{1,j} = -\mathbf{F}_k^{(j)}\mathbf{y}_{2,j}$$

$\quad$to generate approximate solution $\mathbf{x}_{m,j} = \mathbf{x}_{-1,j} + \mathbf{U}_j\mathbf{y}_{1,j} + \mathbf{V}_{m-k}^{(j)}\mathbf{y}_{2,j}$
12: **end for**

Since the matrices are known a priori and only the shifts are varying, we are able to do something different: once the basis $\mathbf{U}$ and $\mathbf{C}$ are known, our approach to update $\mathbf{U}_j$ and $\mathbf{C}_j$ can be easily performed independently and can be parallelized in a straightforward manner. We will discuss the steps and cost for obtaining $\mathbf{C}_j$ at the end of this section. First, let us assume that $\mathbf{U}_j, \mathbf{C}_j$ are available, and discuss our solution technique.

To begin, we assume we have computed $\tilde{\mathbf{x}}_{0,j}$ iteratively and simultaneously for all j using the shift-invariant property as the approximate solution to $(\mathbf{K} + \sigma_j\mathbf{I})\mathbf{x}_j = \mathbf{b}$. We use $\tilde{\mathbf{x}}_{0,j}$ as a first estimate of the solution to $(\mathbf{K} + \sigma_j\mathbf{I} + \sigma'_j\mathbf{R})\mathbf{x} = \mathbf{b}$ obtained by exploiting the shift-invariance of the Krylov subspaces (see Step 1 in Algorithm 1). The corresponding initial residual is denoted $\mathbf{r}_{0,j} := \mathbf{b} - \mathbf{A}_j\tilde{\mathbf{x}}_{0,j}$. Because we expect this solution to need augmentation, we next search for a better estimate of the form $\mathbf{x}_{-1,j} = \tilde{\mathbf{x}}_{0,j} + \mathbf{U}_j\mathbf{z}$. We choose $\mathbf{z}$ such that

$$\mathbf{z} = \arg\min \|\mathbf{b} - \mathbf{A}_j(\tilde{\mathbf{x}}_{0,j} + \mathbf{U}_j\mathbf{z})\| = \arg\min \|\mathbf{r}_{0,j} - \mathbf{C}_j\mathbf{z}\|$$

For this choice, we get a new solution estimate $\mathbf{x}_{-1,j} = \tilde{\mathbf{x}}_{0,j} + \mathbf{U}_j\mathbf{C}_j^T(\mathbf{b} - \mathbf{A}_j\tilde{\mathbf{x}}_{0,j})$ that gives the residual $\mathbf{r}_{-1,j} = (\mathbf{I} - \mathbf{C}_j\mathbf{C}_j^T)\mathbf{r}_{0,j}$.

The approximate solution $\mathbf{x}_{m,j}$ for the parametric system j is obtained by searching in the augmented affine subspace

$$\mathbf{x}_{m,j} \in \mathbf{x}_{-1,j} + \text{Span}\{\mathbf{U}_j\} \oplus \mathcal{K}_{m-k}\left\{(\mathbf{I} - \mathbf{C}_j\mathbf{C}_j^T)\mathbf{A}_j, (\mathbf{I} - \mathbf{C}_j\mathbf{C}_j^T)\mathbf{r}_{0,j}\right\} \tag{3.3}$$

that is, by searching for solutions in the deflated subspace $\mathcal{K}_{m-k}\left\{(\mathbf{I} - \mathbf{C}_j\mathbf{C}_j^T)\mathbf{A}_j, (\mathbf{I} - \mathbf{C}_j\mathbf{C}_j^T)\mathbf{r}_{0,j}\right\}$ obtained by applying $m-k$ steps of the Arnoldi algorithm with the deflated matrix $(\mathbf{I} - \mathbf{C}_j\mathbf{C}_j^T)\mathbf{A}_j$ to the starting vector $(\mathbf{I} - \mathbf{C}_j\mathbf{C}_j^T)\mathbf{b}$ and augmented with the subspace $\mathbf{U}_j$. The initial guess $\mathbf{x}_{-1,j}$ is obtained as $\mathbf{U}_j\mathbf{C}_j^T(\mathbf{b} - \mathbf{A}\tilde{\mathbf{x}}_{0,j})$ and the initial residual is $\mathbf{r}_{0,j} = \mathbf{C}_j\mathbf{C}_j^T\mathbf{b}$, (where $\tilde{\mathbf{x}}_{0,j}$ is the approximation solution obtained by using the shift-invariant property to solve $(\mathbf{K} + \sigma_j\mathbf{I})\mathbf{x}_j = \mathbf{b}$).

The following discussion closely mirrors the presentation of recycling in [34, 39, 23, 33]. The Arnoldi algorithm on the deflated problem yields the matrix relationship

$$(\mathbf{I} - \mathbf{C}_j\mathbf{C}_j^T)\mathbf{A}_j\mathbf{V}_{m-k}^{(j)} = \mathbf{V}_{m-k+1}^{(j)}\bar{\mathbf{T}}_{m-k}^{(j)}$$

where the superscripts indicate the system index j. The above equation can be rewritten as a modified Arnoldi relationship by defining $\mathbf{F}_k^{(j)} \stackrel{\text{def}}{=} \mathbf{C}_j^T\mathbf{A}_j\mathbf{V}_{m-k}^{(j)}$ and reorganizing as

$$\mathbf{A}[\mathbf{U}_j, \mathbf{V}_{m-k}^{(j)}] = [\mathbf{C}_j, \mathbf{V}_{m-k+1}^{(j)}] \begin{bmatrix} \mathbf{I}_k & \mathbf{F}_k^{(j)} \\ \mathbf{0} & \bar{\mathbf{T}}_{m-k}^{(j)} \end{bmatrix} \tag{3.4}$$

Now, $\mathbf{V}_{m-k}^{(j)}$ forms a basis for the subspace $\mathcal{K}_{m-k}\left\{(\mathbf{I} - \mathbf{C}_j\mathbf{C}_j^T)\mathbf{A}_j, (\mathbf{I} - \mathbf{C}_j\mathbf{C}_j^T)\mathbf{r}_{0,j}\right\}$ and now we search for solutions of the form $\mathbf{x}_{m,j} \in \mathbf{x}_{-1,j} + \text{Span}\{\mathbf{U}_j\} \oplus \text{Span}\{\mathbf{V}_{m-k}^{(j)}\}$ and can be written as

$$\mathbf{x}_{m,j} = \mathbf{x}_{-1,j} + \mathbf{U}_j\mathbf{y}_{1,j} + \mathbf{V}_{m-k}^{(j)}\mathbf{y}_{2,j}$$

The solution to the coefficients $\mathbf{y}_{1,j}$ and $\mathbf{y}_{2,j}$ are obtained by minimizing the residual which results in the following least squares problem

$$\min_{\mathbf{y}_{1,j},\mathbf{y}_{2,j}} \left\| \begin{bmatrix} \mathbf{0} \\ \beta_j\mathbf{e}_1 \end{bmatrix} - \begin{bmatrix} \mathbf{I}_k & \mathbf{F}_k^{(j)} \\ \mathbf{0} & \bar{\mathbf{T}}_{m-k}^{(j)} \end{bmatrix} \begin{bmatrix} \mathbf{y}_{1,j} \\ \mathbf{y}_{2,j} \end{bmatrix} \right\| \tag{3.5}$$

where $\beta_j = \|(\mathbf{I} - \mathbf{C}_j\mathbf{C}_j^T)\mathbf{r}_{0,j}\|_2$. Since the number of iterations are expected to be small, we store the vectors $\mathbf{V}_{m-k}^{(j)}$ and solve the least squares problem in equation (3.5) directly. If the number of iterations are expected to be large, we do not need the solution to $\mathbf{y}_{1,j}$ and $\mathbf{y}_{2,j}$ explicitly, only the products $\mathbf{V}_{m-k}^{(j)}\mathbf{y}_{2,j}$ and $\mathbf{F}_k^{(j)}\mathbf{y}_{2,j}$ which can be obtained using short-term recurrence relation similar to recycled MINRES [39] (assuming the matrices $\mathbf{K}$, $\mathbf{M}$ and $\mathbf{R}$ are symmetric).

3.2. Computational and storage costs. We now discuss the computational and storage costs involving Algorithm 1 and the overhead induced by the augmented approach for constructing the augmented basis $\mathbf{U}$. We note that the loop can be executed in parallel because no information is shared across the solves except for the initial choice of $\mathbf{U}$. We focus only on costs that are linear in the size of the matrix, i.e. costs of the form $\mathcal{O}(N)$ since the costs involving smaller matrices are negligible. The augmented approach in Algorithm 1 requires additional storage of $2Nk$ for the matrices $\mathbf{U}$ and $\mathbf{C}$. For the pre-computation of the new bases $\mathbf{U}_j$ and $\mathbf{C}_j$ the major cost is the QR factorization which is $\mathcal{O}(Nk^2)$ for each wavelength. To accelerate this computation, we use a more efficient approach at the possible expense of some accuracy. We first compute

$$\begin{aligned}(\mathbf{C}_j')^T\mathbf{C}_j' = \quad & \mathbf{I} + \sigma_j(\mathbf{C}^T\mathbf{U} + \mathbf{U}^T\mathbf{C}) + \sigma_j'(\mathbf{C}^T\mathbf{R}_\mathbf{U} + \mathbf{R}_\mathbf{U}^T\mathbf{C}) \\ & + \sigma_j\sigma_j'(\mathbf{R}_\mathbf{U}^T\mathbf{U} + \mathbf{U}^T\mathbf{R}_\mathbf{U}) + \sigma_j^2\mathbf{U}^T\mathbf{U} + (\sigma_j')^2\mathbf{R}_\mathbf{U}^T\mathbf{R}_\mathbf{U}.\end{aligned}$$

Next, the small $k \times k$ matrices such as $\mathbf{U}^T\mathbf{C}$ are precomputed and stored since they are independent of the shifts σ_j and σ_j'. Then the Cholesky factorization of $(\mathbf{C}_j')^T\mathbf{C}_j'$ is computed and $\mathbf{C}_j$ is obtained as $\mathbf{C}_j = \mathbf{C}_j'\mathbf{F}^{-1}$. Note as before that the inverse $\mathbf{F}^{-1}$ is not computed explicitly. Because of this pre-computation, the additional cost per wavelength is now only $\mathcal{O}(k^3)$. Suppose the algorithm converges in m iterations, then the algorithm requires $m - k$ additional matrix-vector products and $\mathcal{O}(Nm^2 + mkN)$ other floating point operations. For an efficient algorithm, the overhead costs must be offset by the gains obtained by decreasing the number of iterations. Numerical experiments performed in Section 6 demonstrate that the reduction in the number of matrix-vector products due to deflation offsets the additional computational cost due to pre-computation and re-orthogonalization.

4. Fast compression of H. After computing the incident field $\boldsymbol{\phi}_i$ and the adjoint field $\boldsymbol{\phi}_d$ the next step is to compute the measurement operator $\mathbf{H}$. The construction of $\mathbf{H}$ is described in Equation (2.12). However, recall that the operator $\mathbf{H}$ is expensive to store and compute, and therefore our goal is to produce an approximate factorization of $\mathbf{H} \approx \mathbf{U}\mathbf{V}^T$. Furthermore, an optimal low-rank compression using SVD scales as $\mathcal{O}\left(\min\{M,N\}^2, \max\{M,N\}\right)$ which is prohibitively expensive. Here we present an algorithm that avoids computing $\mathbf{H}$ in its entirety, but computes and compresses sub-blocks and then combines the compressed sub-blocks in a recursive manner.

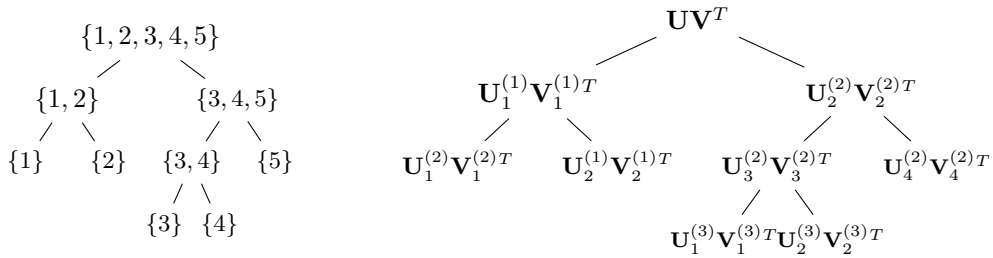

Fig. 4.1: The size of the matrices at the leaf level are $N_{ds}N_\lambda \times N$ and $N_s = 5$. The superscripts denote the level of the tree, where as the subscripts denote the order at a given level (which can be different than the global ordering of the sources). At the leaves the matrices are approximated by low-rank factors and then agglomerated recursively based on the tree. Here we consider a tree corresponding to 5 sources. A possible source configuration that produces the tree can be obtained in Figure 4.2.

4.1. Outline. The rows of $\mathbf{H}$ can be partitioned as

$$\mathbf{H}^T = [\mathbf{H}_1^T, \ldots, \mathbf{H}_{N_s}^T]$$

where each block $\mathbf{H}_i$ is of size $N_{ds}N_\lambda \times N$. Each row of $\mathbf{H}$ represents a discretized version of the integral described in Equation (2.6) that combines the incident field and the adjoint field. Instead of compressing $\mathbf{H}$ by using low-rank factorization techniques such as truncated SVD, the idea is to compress each block $\mathbf{H}_i \approx \mathbf{U}_i\mathbf{V}_i^T$ locally and then combine the factorizations in a recursive fashion. The full algorithm is provided in Algorithm 2 and is illustrated for $N_s = 4$ in Figure 4.1. The low-rank approximation of the blocks $\mathbf{H}_i$ can be accomplished either using Randomized SVD or partially pivoted ACA algorithms described in Subsection 4.2. Then a scheme for ordering the blocks $\mathbf{H}_i$ is presented in Subsection 4.4 that uses a spatial bisection tree to order the source locations by their spatial proximity. The agglomeration of the low-rank factors is accomplished by recursion using this tree structure. In Subsection 4.5 we analyze the computational costs of this recursive compression scheme and conclude with an error analysis in Subsection 4.6.

4.2. Low Rank representation. As was mentioned earlier, we require a strategy to compute a low-rank factorization of the sub-blocks $\mathbf{H}_j^T$ for $j = 1, \ldots, N_s$. Consider the blocks $\mathbf{H}_j$ of size $m \times n$ where $m = N_{ds}N_\lambda$ and $n = N$. Considering only one source, we seek a low-rank approximation of the form $\mathbf{H} \approx \mathbf{U}\mathbf{V}^T$ and the number of columns of $\mathbf{U}$ and $\mathbf{V}$ is denoted by r. It is well known that the best low-rank approximation to rank r is obtained by truncating the SVD to rank r. In this case, we have $\|\mathbf{H} - \mathbf{U}_r\Sigma_r\mathbf{V}_r^T\|_2 = \sigma_{r+1}$. However, computing the SVD is expensive since it requires $\mathcal{O}(nm^2)$ operations assuming $n \geq m$. The advantage of the low-rank representation is that the cost of storing the decomposition and computing matrix-vector products are both given by $\mathcal{O}(r(m+n))$ instead of $\mathcal{O}(mn)$. When $r \ll \min\{m, n\}$, this reduction can represent significant savings.

In the Appendix A we describe two approximate methods that compute a low-rank representation but have a lower computational cost asymptotically than the SVD, namely randomized SVD (RandSVD) and partially pivoted Adaptive Cross Approximation (ppACA). Here, we only summarize the resulting computational costs. For a matrix of dimensions $m \times n$ that has a rank r the cost can be expressed as

$$\mu_{\text{Comp}}(m, n; r) \quad = \quad \begin{cases} C_1 rmn + C_2 r^2(m+n) & \text{RandSVD} \\ C_3 r^2(m+n) & \text{ppACA} \end{cases} \tag{4.1}$$

where constants C_1, C_2 and C_3 are assumed to be known and provided in the literature (see Appendix A). Note that the costs of the low-rank approximation both using RandSVD and ppACA are asymptotically smaller than the cost of an SVD which scales as $\mathcal{O}(nm^2)$ assuming $n > m$.

4.3. Agglomerating low ranks. Having produced low-rank approximations to the sub-blocks $\mathbf{H}_j$ for $j = 1, \ldots, N_s$ we now consider the problem of agglomerating low-rank factors to produce a global low-rank factorization. Here we consider only two sources, i.e. $N_s = 2$ and as before, $\mathbf{H}^T = [\mathbf{H}_1^T, \mathbf{H}_2^T]$ of size $m \times n$ where $m = N_{ds} N_\lambda$ and $n = N$. Suppose we have the low-rank factors $\mathbf{H}_1 \approx \mathbf{U}_1 \mathbf{V}_1^T$ and $\mathbf{H}_2 \approx \mathbf{U}_2 \mathbf{V}_2^T$ each with rank r which have been compressed according to some predetermined tolerance ε. The low-rank factors can be then combined as

$$\mathbf{H} = \begin{bmatrix} -\mathbf{H}_1 - \\ -\mathbf{H}_2 - \end{bmatrix} \approx \begin{bmatrix} \mathbf{U}_1 \mathbf{V}_1^T \\ \mathbf{U}_2 \mathbf{V}_2^T \end{bmatrix} = \begin{bmatrix} \mathbf{U}_1 & \\ & \mathbf{U}_2 \end{bmatrix} \begin{bmatrix} -\mathbf{V}_1^T - \\ -\mathbf{V}_2^T - \end{bmatrix}.$$

Now the leftmost matrix is $m \times r$ and has independent columns by construction. The rightmost $r \times n$ matrix, however, may have a rank smaller yet than $\min r, n$. So we compute $\begin{bmatrix} -\mathbf{V}_1^T - \\ -\mathbf{V}_2^T - \end{bmatrix} = \mathbf{U}_\mathbf{V} \mathbf{V}^T$ where $\mathbf{U}_\mathbf{V}, \mathbf{V}$ have $r' < r$ columns, and we set

$$\begin{bmatrix} \mathbf{U}_1 & \\ & \mathbf{U}_2 \end{bmatrix} \mathbf{U}_\mathbf{V} \mathbf{V}^T = \mathbf{U} \mathbf{V}^T, \qquad \mathbf{U} \in \mathbb{R}^{2N_{ds} N_\lambda \times r'}, \mathbf{V} \in \mathbb{R}^{N \times r'}.$$

We will need to form $\mathbf{U}$ explicitly, and this requires computing a rank-r' approximation to the stacked $\mathbf{V}_i$ matrix. The cost must therefore include the matrix-matrix product (using the block diagonal structure of the leftmost matrix in the product) and the cost of this additional factorization, using the same algorithm as was used to approximately factor the $\mathbf{H}_i$. The estimate of the cost of agglomeration of sub-blocks is therefore,

$$\mu_{\text{Agg}}(m, n; r', r) \quad = \quad \begin{cases} C' r r' (m+n) & \text{RandSVD} \\ C'' (r')^2 (m+n) & \text{ppACA} \end{cases} \tag{4.2}$$

where the constant C' and C'' are assumed to be known (see Appendix A). For simplicity of analysis we will only use RandSVD for agglomerating low-rank representations.

4.4. Recursive SVD. We have, thus far, described a procedure to compute low-rank factors $\mathbf{H}_j \approx \mathbf{U}_j \mathbf{V}_j^T$ corresponding to source indexes $j = 1, \ldots, N_s$ and shown how to agglomerate low-rank factors when the number of sources are 2. However, as we shall show, this procedure can be implemented recursively. The complete algorithm for computing and compressing $\mathbf{H}$ is summarized in Algorithm 2.

Let $I = \{1, \ldots, N_s\}$ denote an index set and let $|I|$ denote the cardinality of the set I. Let the index sets I be partitioned into binary trees denoted by T_I respectively. For each $\tau \in T_I$, we denote the set of its sons by $S(\tau) \subset T_I$ and the leaves of the tree are denoted by $\mathcal{L}(T_I)$. The tree T_I can be used to recursively spatially partition the domain so that the sources are ordered such that recursively combining low-rank factors from sub-blocks corresponds to combining nearby sources since the corresponding Green's functions are known to be highly compressible [9]. As a result, by this ordering, we hope to gain a larger compression at each level in the tree. Given the tree T_I we can recursively compress the sub-blocks using Algorithm 2 which is initialized using the root of the tree, which corresponds to the index set I. The partitioning of the matrix $\mathbf{H}$ into sub-blocks and their recursive compression for $N_s = 4$ using the Figure 4.1. The construction of the tree is described in Algorithm 2.

We make the assumption that the locations of the source, detectors and the support of the perturbation are non-overlapping and well separated. For concreteness, let us assume that the breast is placed between parallel plates and is enclosed by a cuboid of dimensions $[-L_x, L_x] \times [-L_y, L_y] \times [0, L_z]$. Furthermore, let us assume that the sources are located on the top plate $z = L_z$ and the detectors are located at $z = 0$ and the detectors for any given source are roughly co-axial. The construction of the tree is performed as follows, and a simple example illustrating this construction is provided in Figure 4.2. Let J be the index set corresponding to source locations. The tree T_J is constructed recursively by geometric bisection applied on a 2D domain. Given the initial bounding box containing all the source (in this case, a box of size $[-L_x, L_x] \times [L_y, L_y]$), the domain is split in a direction perpendicular to the direction of maximum expanse and the sources are split between the newly

Algorithm 2 RecursiveLowRank(τ, ε)

if $S(\tau) \neq \emptyset$ **then**
  $\mathcal{U} := \emptyset$, $\mathcal{V} := \emptyset$
  for all $\tau' \in S(\tau)$ **do**
    $\mathbf{U}_{\tau'}, \mathbf{V}_{\tau'} = \text{RecursiveLowRank}(\tau', \varepsilon)$
    $\mathcal{U}$.append($\mathbf{U}$), $\mathcal{V}$.append($\mathbf{V}$)
  end for
else
  for all $\tau' \in \mathcal{L}(T_I)$ **do**
    Compute the sub-block $\mathbf{H}'_\tau$ using Equation (2.6)
    {// Use fast solvers developed in Section 3 for incident $\boldsymbol{\phi}_i$ and adjoint field $\boldsymbol{\phi}_d$. }
    $[\mathbf{U}, \mathbf{V}] = \text{LowRank}(\mathbf{H}|_{\tau'}, \varepsilon)$
    {//Compute low-rank at leaf level such that $\|\mathbf{H}_{\tau'} - \mathbf{U}\mathbf{V}^T\| \leq \varepsilon \|\mathbf{H}_{\tau'}\|$. see Section 4.2}
    $\mathcal{U}$.append($\mathbf{U}$), $\mathcal{V}$.append($\mathbf{V}$)
  end for
end if
$[\mathbf{U}, \mathbf{V}] = \text{Agglomerate}(\mathcal{U}, \mathcal{V}, \varepsilon)$ {// Agglomerate low-rank factors, see Section 4.3}
return $\mathbf{U}, \mathbf{V}$ such that $\mathbf{H} \approx \mathbf{U}\mathbf{V}^T$

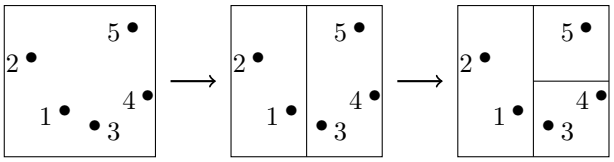

Fig. 4.2: Top view of the bounding box for a possible source configuration. Locations of 5 different sources that can be subdivided recursively using Algorithm 3. The resulting tree T_J that is returned by the algorithm, when initialized by the index set $J = \{1, \ldots, 5\}$ is provided in Figure 4.1.

created split domains. If the number of points in each domain are greater than 2, the procedure is computed recursively until the number of points in each leaf is no greater than 2. If the points are uniformly distributed, the number of levels $L = \log_2(N_s)$, where N_s is the number of sources. The algorithm is summarized in Algorithm 3. The algorithm is initialized using $\tau = J$, $\alpha = [-L_x, L_x]$ and $[-L_y, L_y]$.

Algorithm 3 Split(τ, α, β) // Geometric bisection to construct tree

$j_{\max} := \arg\max\{\beta_j - \alpha_j : j \in \{1, \ldots, d\}\}$
$\gamma := (\alpha_{j_{\max}} + \beta_{j_{\max}})/2$ {// Split cluster in chosen direction }
$\tau_1 := \emptyset$, $\tau_2 := \emptyset$
for $i \in \tau$ **do**
  if $e_{j_{\max}}^T \mathbf{x}_i \leq \gamma$ **then**
    $\tau_1 = \tau_1 \bigcup \{i\}$
  else
    $\tau_2 = \tau_2 \bigcup \{i\}$
  end if
end for
Define $\beta' := \beta$, $\beta'_{j_{\max}} := \gamma$, and $\alpha' := \alpha$, $\alpha'_{j_{\max}} := \gamma$
Split(τ_1, α, β') and Split(τ_2, α', β) {// Split only if $|\tau_i| > 2$ for $i = 1, 2$.}

4.5. Computational cost. In this section we present a detailed analysis of the computational cost of the recursive SVD described in Algorithm 2. Our analysis is in the same spirit as the one

provided in [18]. However, while they used the full SVD algorithm to produce the optimal truncation at each level, we will consider the use of approximate low-rank factorizations that have better computational costs. As we shall see, this analysis and results are more sharp than [18] since we account for the cost depending on rank of the sub-blocks at each level of the tree.

The starting point of our analysis is splitting the costs into a contribution that comes from the preparation at the leaf level that is computed using the low-rank factorization techniques listed in Section 4.2 and the contribution that comes from agglomerating low-rank sub-blocks. The number of sub-blocks to agglomerate become fewer the higher up we are in the tree, but the size of the matrices to be agglomerated increases. The total cost is therefore equal to

$$\text{Cost} \;=\; \sum_{b\in\mathcal{L}(T_I)} \mu_{Comp}(m^{(L)}, n^{(L)}; r_L) + \sum_{\ell=0}^{L-1} \sum_{b\in T_I^{(l)}} \mu_{\text{Agg}}\left(m^{(\ell)}, n^{(\ell)}; r_{L-\ell-1}, r_{L-\ell}\right) \text{ flops.} \tag{4.3}$$

We assume that N_s is a power of 2 and each partition has the same size, so that the size of the partition at the leaf level is $N_\lambda N_{ds} \times N$, and $N_s = 2^L$ where L is the depth of the tree T_I. Furthermore, we denote by r_ℓ, the maximum block rank of any partition at level ℓ. The cost of recursive SVD is then

$$\begin{aligned}
\text{Cost} &= 2^L \mu_{\text{Comp}}\left(\frac{M}{2^L}, N; r_L\right) &+ \sum_{\ell=0}^{L-1} 2^{L-\ell-1}\mu_{\text{Agg}}\left(2^{\ell-L}M, N; r_{L-\ell-1}, r_{L-\ell}\right)\\
&= C_1 2^L r_L \frac{M}{2^L} N + 2^L C_2 r_L^2\left(\frac{M}{2^L}+N\right) &+ \sum_{\ell=0}^{L-1} C' 2^{L-\ell-1} r_{L-\ell-1} r_{L-\ell}(2^{\ell-L}M+N)\\
&= C_1 r_L MN + C_2 r_L^2\left(M + 2^L N\right) &+ \sum_{\ell=0}^{L-1} C' 2^{L-\ell-1} 2\delta_{L-\ell}^{L-\ell-1} r_{L-\ell}^2(2^{\ell-L}M+N)\\
&= C_1 r_L MN + C_2 r_L^2\left(M + 2^L N\right) &+ 2^{L-1}C' r_L^2 \sum_{\ell=0}^{L-1}(M/2^L + 2^{-\ell}N) 2\delta_{L-\ell}^{L-\ell-1} \prod_{k=L-\ell}^{L-1}\left(2\delta_{k+1}^k\right)^2
\end{aligned}$$

Here, we assume that $r_\ell = 2\delta_{\ell+1}^\ell r_{\ell+1}$, that is, the ratio of the ranks at level ℓ that is obtained by agglomeration of two sub-blocks each of rank at most $r_{\ell+1}$ is a factor of $\delta_{\ell+1}^\ell$, which is dependent on the particular level under consideration. It is easy to see that $0 < \delta_{\ell+1}^\ell \leq 1$ for all $\ell = 0, \ldots, L-1$. If we further make the assumption that $\max_{0\leq\ell<L} \delta_{\ell+1}^\ell = \delta$ is independent of the level ℓ, then we have following result that the total computational cost of recursive low-rank truncation can be further simplified to

$$C_1 r_L MN + C_2 r_L^2\left(M + 2^L N\right) + 2^L C' r_L^2 \delta\left(\frac{1-(\sqrt{2}\delta)^{2L-1}}{1-\sqrt{2}\delta}\frac{M}{2^L} + \frac{1-(2\delta)^{2L-1}}{1-2\delta}N\right). \tag{4.4}$$

We define the quantities f_1 and f_2 as follows:

$$\begin{aligned}
f_1 \stackrel{\text{def}}{=} \sum_{\ell=0}^{L-1} 2^{-\ell}\delta_{L-\ell}^{L-\ell-1} \prod_{k=L-\ell}^{L-1}\left(2\delta_{k+1}^k\right) &\leq \delta\frac{1-(\sqrt{2}\delta)^{2L-1}}{1-\sqrt{2}\delta}\\
f_2 \stackrel{\text{def}}{=} \sum_{\ell=0}^{L-1} \delta_{L-\ell}^{L-\ell-1} \prod_{k=L-\ell}^{L-1}\left(2\delta_{k+1}^k\right) &\leq \delta\frac{1-(2\delta)^{2L-1}}{1-2\delta}.
\end{aligned}$$

Here, we assume that $\delta \neq 1/\sqrt{2}, 1/2$. If this were the case, the appropriate sums would simplify to equal L. In the worst case, when there is no compression at higher levels $\delta = 1$ and then $f_1 \sim 2^L$ and $f_2 \sim 4^L$.

We now compare the computational costs between different methods that we have outlined in this section. It can be readily seen that the cost of the recursive factorizations scale asymptotically

Method	Leaf computation		Tree computation	
	RandSVD	ppACA	Average	Worst
Direct	$C_1RMN + C_2R^2(M+N)$	$C_3R^2(M+N)$	-	-
Recursive	$C_1r_LMN + C_2r_L^2(M+2^LN)$	$C_3r_L^2(M+2^LN)$	$Cr_L^2(Mf_2+2^LNf_1)$	$r_L^2C4^L(M+N)$

Table 4.1: Summary of computational costs of the recursive SVD algorithm. Here $M = N_sN_{ds}N_\lambda$ is the number of measurements and N is the grid size. The storage costs are $\mathcal{O}(R(M+N))$ where R is the global rank of the low-rank factorization.

better than the SVD which scales as $\mathcal{O}(NM^2)$ assuming $M \leq N$. We denote by 'Direct', the low-rank algorithms described in Section 4.2. Considering only the costs that are of the order $\mathcal{O}(MN)$, it can be readily observed that simply the recursive SVD methods have better scaling than the 'Direct' method which scales as $\mathcal{O}(NM^2)$, if the rank at the leaves are smaller than the overall rank of the matrix $r_L < R$. By comparing the worst case costs for the recursive SVD methods, we can see that that the terms linear in M and N are comparable if $r_L \sim R/2^L$. If no compression is observed at any levels, including the leaves, then our algorithm performs poorly since we are needlessly computing a large number of "low-rank" factorizations at all the levels in the tree. However, there are still a couple of benefits of using the recursive SVD approach. First, for the range of parameters we are interested in exploring, storing the entire matrix $\mathbf{H}$ could cost ~ 200 GB which may be completely infeasible to store and later compress. By contrast the strategy in Algorithm 2 does not require storage of $\mathbf{H}$ in its entirety but computes and compresses sub-blocks of $\mathbf{H}$ on-the-fly and therefore, has favorable storage costs. Second, our algorithm provides more locality in the calculations and therefore, the algorithm is more amenable to parallelization and distributed computing setting. This has also been noted by [9].

To summarize the asymptotic cost of factorization using RandSVD at the leaf level is $\mathcal{O}(r_LMN + r_L^2(N+M))$ and using ppACA it is $\mathcal{O}(r_L^2(N+M))$. Numerical evidence suggests that there is compression at every level (and therefore $r_L \ll R$) and this justifies the use of this hierarchical approach.

4.6. Accuracy. We now discuss the accuracy of the recursive low-rank approximation. In the algorithms described above, there are two sources of error - due to the low-rank truncation at the leaf level and the error accumulated due to the agglomeration process at all other levels in the tree. In order to analyze the accuracy of the recursive SVD computation, we first consider the accuracy of the agglomeration step at one level. We consider the matrix $\mathbf{H}$ which has a partitioned as $\mathbf{H} = [\mathbf{H}_1^T, \mathbf{H}_2^T]^T$. Suppose we compute a low-rank approximation to $\mathbf{H}_i \approx \hat{\mathbf{H}}_i \stackrel{\text{def}}{=} \mathbf{U}_i\mathbf{V}_i^T$ for $i = 1, 2$ using the techniques described in Section 4.2. We assume that the low-rank matrices satisfy the bounds $\|\mathbf{H}_i - \hat{\mathbf{H}}_i\| \leq \varepsilon\|\mathbf{H}_i\|$ for $i = 1, 2$. We can then bound the error in the approximation $\mathbf{H} \approx \hat{\mathbf{H}} \stackrel{\text{def}}{=} \mathbf{U}\mathbf{V}^T$.

$$\begin{aligned}
\|\mathbf{H} - \hat{\mathbf{H}}\| &\leq \|\hat{\mathbf{H}} - [\hat{\mathbf{H}}_1^T, \hat{\mathbf{H}}_2^T]\| + \|\mathbf{H}_1 - \hat{\mathbf{H}}_1\| + \|\mathbf{H}_2 - \hat{\mathbf{H}}_2\| \\
&\leq \varepsilon\left(\|[\hat{\mathbf{H}}_1^T, \hat{\mathbf{H}}_2^T]^T\|\right) + \varepsilon\left(\|\mathbf{H}_1\| + \|\mathbf{H}_2\|\right) \\
&\leq \varepsilon\left(\|\hat{\mathbf{H}}_1\| + \|\hat{\mathbf{H}}_2\|\right) + \varepsilon\left(\|\mathbf{H}_1\| + \|\mathbf{H}_2\|\right) \\
&\leq (2\varepsilon + \varepsilon^2)\left(\|\mathbf{H}_1\| + \|\mathbf{H}_2\|\right) = 2\varepsilon\left(\|\mathbf{H}_1\| + \|\mathbf{H}_2\|\right) + \mathcal{O}(\varepsilon^2)
\end{aligned}$$

We have used the fact that the strategy that is used to truncate the rank for the agglomeration is the same as one to compute the low-rank compression. Furthermore, we have also have used the inequality that

$$\|\hat{\mathbf{H}}_i\| = \|\hat{\mathbf{H}}_i - \mathbf{H}_i + \mathbf{H}_i\| \leq \|\hat{\mathbf{H}}_i - \mathbf{H}_i\| + \|\mathbf{H}_i\| \leq (1+\varepsilon)\|\mathbf{H}_i\|$$

We now extend it to the case where $N_s > 2$ by recursively applying the error bound that was derived above

$$\begin{aligned}\|\mathbf{H} - \hat{\mathbf{H}}\| &\leq \sum_{b\in\mathcal{L}(T_I)} \|\mathbf{H}_b - \hat{\mathbf{H}}_b\| + \sum_{\ell=0}^{L-1} \sum_{b\in T_I^{(l)}} \|\hat{\mathbf{H}}_b - \text{Agg}\{\hat{\mathbf{H}}_{b'} : b' \in S(b))\}\| \\ &\leq (L+1)\varepsilon \left(\sum_{b\in\mathcal{L}(T_I)} \|\mathbf{H}_b\| \right) + \mathcal{O}(\varepsilon^2)\end{aligned}$$

If the computations were performed in the Frobenius norm, then using the Cauchy-Schwarz inequality we can conclude that $\|\mathbf{H} - \hat{\mathbf{H}}\|_F \lesssim 2^{L/2}(L+1)\varepsilon\|\mathbf{H}\|_F$. In order to derive an equivalent relationship for the 2-norm, we use the following inequality

$$\|\mathbf{H} - \hat{\mathbf{H}}\|_2 \leq \|\mathbf{H} - \hat{\mathbf{H}}\|_F \lesssim 2^{L/2}(L+1)\varepsilon\|\mathbf{H}\|_F \leq \sqrt{N_r}2^{L/2}(L+1)\varepsilon\|\mathbf{H}\|_2$$

where $N_r = \min\{M, N\}$. Therefore, in order to achieve a desired relative tolerance ε_d, the tolerance that is used in the low-rank approximation and the agglomeration can be computed as $\varepsilon \sim \varepsilon_d/2^{L/2}(L+1)\sqrt{N_r}$ for the Recursive SVD.

This error bound although locally optimal can result in a low-rank factorization that may be suboptimal in terms of compression. For this reason, we propose an additional step for compressing the low-rank factors $\hat{\mathbf{H}} = \mathbf{U}\mathbf{V}^T$.

1. Compute thin QR factorizations $\mathbf{Q_U}\mathbf{R_U} = \mathbf{U}$ and $\mathbf{Q_V}\mathbf{R_V} = \mathbf{V}$
2. Compute SVD $\mathbf{R_U}\mathbf{R_V}^T = \mathbf{U}'\mathbf{\Sigma}'(\mathbf{V}')^T$
3. Truncate R singular values and return $\mathbf{U} = \mathbf{Q_U}\mathbf{U}'_r$ and $\mathbf{V} = \mathbf{Q_V}\mathbf{V}'_r\mathbf{\Sigma}'_r$

This additional cost is $\mathcal{O}(R^2(M+N) + R^3)$ and may be beneficial when $R \ll \min\{M, N\}$.

The take away is that the cost of storage and matvecs with $\hat{\mathbf{H}}$ is $\mathcal{O}(R(M+N))$, which is critical when we need to access and multiply with the estimate repeatedly in the course of the optimization for the image parameters. We now describe that optimization problem.

5. Reconstruction algorithms. The recovery of the shape of the tumor and the chromophore concentrations from diffuse optics measurements is an ill-posed inverse problem. The inverse problem can be stated as follows: Given a set of measurements $\mathbf{y}$ that measures the scattered field $\boldsymbol{\phi}_s$ at multiple detector locations $\mathbf{r}_d$ due to incident field $\boldsymbol{\phi}_i$ from multiple source locations and illuminated at several different wavelengths, recover the spatially varying perturbation of absorption $\Delta\mu_a(\mathbf{r}, \lambda)$ and the concentration of the chromophore species. Standard approaches to deal with ill-posedness introduce some kind of regularization, such as Tikhonov regularization. Here, we consider the parametric level setup approach proposed in [1] (abbreviated as PaLS) and subsequently applied to the diffuse optical tomography problem in [27].

We briefly review the PaLS approach for parameterizing the shape perturbation. The characteristic function $\chi(\mathbf{r})$ defined in equation (2.8) is taken as the τ-level set of a Lipschitz continuous function $\varphi(\mathbf{r}) : \mathcal{D} \to \mathbb{R}$. Using $\varphi(\mathbf{r})$, the characteristic function $\chi(\mathbf{r})$ can be expressed as

$$\chi(\mathbf{r}) = H\left(\varphi(\mathbf{r}) - \tau\right) \qquad \varphi(\mathbf{r}) = \sum_{k=1}^{n_p} \alpha_k \psi\left(\beta_k \|\mathbf{r} - \boldsymbol{\chi}_k\|^\dagger\right) \tag{5.1}$$

where $H(\cdot)$ is the Heaviside function. In practice, we use smooth approximations H_ε of the Heaviside function H, and its derivative denoted by δ_ε. We represent the function $\varphi(\mathbf{r})$ parametrically as weighted combinations of basis functions $\psi(\cdot)$ and we have $\|\mathbf{r}\|^\dagger = \sqrt{\|\mathbf{r}\|_2^2 + \nu^2}$ and $\nu > 0$ is a small parameter to ensure that φ is differentiable. Several choices are available for ψ such as polynomials and radial basis functions. Here we choose the compactly supported radial basis functions that were previously used in [1]. The coefficients α_k control the magnitude of the radial basis functions, β_k control the width and $\boldsymbol{\chi}_k$ control the centers. The basis functions and their number control how fine or coarse the representation will be. On the one hand, having a large number of basis functions will be beneficial in reconstructing fine scale features, however, it has additional associated computational cost and further exacerbates the non-convexity.

The parameters that need to be estimated are collected in a vector $\mathbf{p} = [\boldsymbol{\alpha}^T, \boldsymbol{\beta}^T, \boldsymbol{\chi}_x^T, \boldsymbol{\chi}_y^T, \boldsymbol{\chi}_z^T]^T \in \mathbb{R}^{N_p}$, where $N_p = (d+2)n_p$ and d is the dimension of the problem. The reconstruction problem can now be stated as the minimization of the following functional

$$\hat{\mathbf{c}}, \hat{\mathbf{p}} \quad \stackrel{\text{def}}{=} \quad \arg\min_{\mathbf{c},\mathbf{p}} \ \|\boldsymbol{\varepsilon}\|_2^2 \ = \ \|\mathbf{W}(\mathbf{y} - \mathbf{D}(\mathbf{p})\mathbf{c})\|_2^2 \tag{5.2}$$

where the columns of $\mathbf{D}(\mathbf{p})$ are given by $\mathbf{E}_i\mathbf{H}\boldsymbol{\mu}(\mathbf{p})$ and $\mathbf{c} = [c_1, \ldots, c_{N_{\text{sp}}}]^T$ represent the concentration of the chromophores.

Algorithm 4 Optimization procedure for solving shape parameters $\mathbf{p}$ and chromophore concentrations $\mathbf{c}$

1: Given tolerances τ_1 and τ_2 and initial guess for PaLS parameters $\mathbf{p}$
2: **while** $\|\boldsymbol{\varepsilon}\|_2 \leq \tau_1$ **do**
3: $\quad \mathbf{c} = (\mathbf{W}\mathbf{D}(\mathbf{p}))^+ (\mathbf{W}\mathbf{y})$
4: $\quad$ **while** $\|\boldsymbol{\varepsilon}\|_2 \leq \tau_2$ **do**
5: $\quad\quad (\mathbf{J}^T\mathbf{J} + \nu\mathbf{I})\delta\mathbf{p} = -\mathbf{J}^T\boldsymbol{\varepsilon}$
$\quad\quad$ {//The parameter ν is chosen by a Levenberg-Marquardt procedure. }
6: $\quad\quad \mathbf{p} \leftarrow \mathbf{p} + \delta\mathbf{p}$
7: $\quad$ **end while**
8: **end while**
9: **return** Shape parameters $\mathbf{p}$ and chromophore concentration $\mathbf{c}$

The resulting optimization problem is solved by alternating between solving for the concentration parameters $\mathbf{c}$ which is a linear least-squares problem and solving for the PaLS parameters $\mathbf{p}$ using a Levenberg-Marquardt procedure. The optimization algorithm requires constructing the Jacobian

$$\mathbf{J} = \frac{\partial \boldsymbol{\varepsilon}}{\partial \mathbf{p}} = -\mathbf{W}\bar{\mathbf{E}}\mathbf{H}\frac{\partial \boldsymbol{\mu}}{\partial \mathbf{p}}$$

where $\bar{\mathbf{E}} = \sum_{i=1}^{N_{\text{sp}}} c_i\mathbf{E}_i$. Analytical expressions for the derivatives $\frac{\partial \boldsymbol{\mu}}{\partial \mathbf{p}}$ are provided in [1]. The stopping criteria for the iterative procedure is chosen according to the discrepancy principle, i.e., the iterations are terminated when the norm of the residuals is less than the noise level up to a user defined constant $\gamma > 1$. In mathematical terms, the stopping criterion becomes $\|\boldsymbol{\varepsilon}\|_2 \leq \gamma\|\boldsymbol{\eta}\|_2$ and $\boldsymbol{\eta}$ is the noise defined in Equation (2.12). More efficient algorithms are available for the reconstruction of PaLS parameters, for example, see TREGS [12]. However, we have chosen the Levenberg-Marquardt algorithm for its relative simplicity of implementation.

Recall in Section 4, we used a compressed low-rank representation of the measurement operator $\mathbf{H} \approx \mathbf{U}\mathbf{V}^T$. Let us denote $\hat{\mathbf{H}} = \mathbf{U}\mathbf{V}^T$ and the error as $\mathbf{E}_\mathbf{H}$ in the low-rank truncation process such that $\mathbf{H} = \hat{\mathbf{H}} + \mathbf{E}_\mathbf{H}$. From the results in Section 4 we know that $\|\mathbf{E}_\mathbf{H}\| \leq \varepsilon\|\mathbf{W}\bar{\mathbf{E}}\|\|\mathbf{H}\|$ and define $\bar{\varepsilon} \stackrel{\text{def}}{=} \varepsilon\|\mathbf{W}\bar{\mathbf{E}}\|$. To simplify the theoretical analysis, we rescale $\mathbf{y} \leftarrow \mathbf{W}$, $\mathbf{H} \leftarrow \mathbf{W}\bar{\mathbf{E}}\mathbf{H}$ and $\hat{\mathbf{H}} \leftarrow \mathbf{W}\bar{\mathbf{E}}\hat{\mathbf{H}}$. The approximate Jacobian $\bar{\mathbf{J}}$ is now given by the expression $\bar{\mathbf{J}} = -\bar{\mathbf{H}}\frac{\partial \boldsymbol{\mu}}{\partial \mathbf{p}}$, so that we have $\mathbf{J} = \bar{\mathbf{J}} + \mathbf{E}_\mathbf{J}$, where $\mathbf{E}_\mathbf{J} = \mathbf{E}_\mathbf{H}\frac{\partial \boldsymbol{\mu}}{\partial \mathbf{p}}$. It can be readily shown that the approximation to the objective function $\mathbf{f}(\mathbf{p}) \stackrel{\text{def}}{=} \|\mathbf{y} - \mathbf{H}\boldsymbol{\mu}(\mathbf{p})\|_2^2$ and the gradient $\nabla_\mathbf{p}\mathbf{f}(\mathbf{p}) \stackrel{\text{def}}{=} -\frac{\partial \boldsymbol{\mu}}{\partial \mathbf{p}}^T\mathbf{H}^T(\mathbf{y} - \mathbf{H}\boldsymbol{\mu}(\mathbf{p}))$ (and the equivalent quantities $\hat{\mathbf{f}}$ and $\nabla_\mathbf{p}\hat{\mathbf{f}}$ with the approximation $\bar{\mathbf{H}}$ instead of $\mathbf{H}$) satisfy the following approximation bounds

$$\begin{aligned} |\mathbf{f} - \hat{\mathbf{f}}| \quad &\leq \quad 2\bar{\varepsilon}\|\mathbf{y} - \hat{\mathbf{H}}\boldsymbol{\mu}\|_2\|\mathbf{H}\|_2\|\boldsymbol{\mu}\|_2 + \mathcal{O}(\bar{\varepsilon}^2) \\ \|\nabla_\mathbf{p}\mathbf{f} - \nabla_\mathbf{p}\hat{\mathbf{f}}\| \quad &\leq \quad \bar{\varepsilon}\|\frac{\partial \boldsymbol{\mu}}{\partial \mathbf{p}}\|_2\|\mathbf{H}\|_2\left(\|\boldsymbol{\mu}\|_2\|\mathbf{H}\|_2 + \|\mathbf{y} - \hat{\mathbf{H}}\boldsymbol{\mu}\|_2\right) \end{aligned}$$

Furthermore, assume that $\|\nabla_\mathbf{p}\mathbf{f} - \nabla_\mathbf{p}\hat{\mathbf{f}}\| \leq \tau_g\|\nabla_\mathbf{p}\hat{\mathbf{f}}\|$, then the acute angle θ between the gradient

$\nabla_{\mathbf{p}}\mathbf{f}$ and $\nabla_{\mathbf{p}}\hat{\mathbf{f}}$ satisfies the following inequality

$$\cos\theta \geq \frac{1-\tau_g^2}{\sqrt{1+\tau_g^2}}$$

The result follows from the result in [40, Lemma 3.1]. We assume that the angle θ is acute, i.e., $\tau_g < 1$ which is always possible since we can control the error $\bar{\varepsilon}$ and therefore the tolerance τ_g. We now present a result that bounds the error between the true and the perturbed search directions in step 5 of Algorithm 4.

PROPOSITION 1. *Assume that the derivative $\frac{\partial \boldsymbol{\mu}}{\partial \mathbf{p}}$ is full rank and $N_p \leq R$, where R is the effective rank of the low-rank representation $\hat{\mathbf{H}}$ and N_p is the number of PaLS parameters. In Algorithm 4 let $\delta\mathbf{p}$ be the search direction corresponding to the exact Jacobian $\mathbf{J}$ and let $\delta\bar{\mathbf{p}}$ be the search direction corresponding to the approximate Jacobian $\bar{\mathbf{J}}$. Then, we can bound the error between the two search directions as*

$$\|\delta\mathbf{p} - \delta\bar{\mathbf{p}}\|_2 \leq \left[\eta(\nu)\|\delta\bar{\mathbf{p}}\|_2 + \frac{\|\boldsymbol{\varepsilon}\|_2}{\nu + \bar{\sigma}_{N_p}^2}\right] \|\mathbf{E}_{\mathbf{J}}\|_2 \tag{5.3}$$

where the factor $\eta(\nu) = \max_{\bar{\sigma}_{N_p} \leq \sigma \leq \bar{\sigma}_1} \sigma/(\nu+\sigma^2)$ and $\bar{\sigma}_1$ and $\bar{\sigma}_{N_p}$ are upper and lower bounds for the singular values of the unperturbed Jacobian $\mathbf{J}$.

The proof is readily obtained by an application of the result [22, Theorem 3.1]. The matrix $\frac{\partial \boldsymbol{\mu}}{\partial \mathbf{p}}$ is full-rank and the dimension of $\mathbf{p}$ denoted by N_p is smaller than the rank of $\hat{\mathbf{H}}$, therefore the Jacobian $\mathbf{J}$ and perturbed Jacobian $\bar{\mathbf{J}}$ are full-rank and satisfy the requirements of the theorem. If $R < N_p$ then the Jacobian is rank-deficient and we could consider a subset selection procedure similar to [22].

6. Numerical Experiments. We present some results of the algorithms that we described in Section 3. For the rest of this section, we consider the following test problem. The geometry under consideration is a breast shaped phantom that is compressed between two flat plates (see Figure 1.1). At its widest it is 12 cm long and the maximum thickness is 5 cm. The domain is discretized using gmsh [13], an open source 3D finite element mesh generator. The finite element matrices corresponding to the discretized representations of the partial differential equations given by equation (2.11) are computed using FEniCS [30] accessed using its Python interface. The boundary $\partial\Omega_R$ for which refractive index mismatch conditions are applied are assumed to be the flat top portions of the boundary, where as zero Dirichlet boundaries are applied on the rest of the boundary.

We consider the background medium to be composed of $N_{\text{sp}} = 4$ species, oxygenated and deoxygenated hemoglobin denoted as HbO_2 and HbR respectively, water H_2O and lipids. These species have been specifically chosen since they are the most optically active chromophores, found in breast tissue, in the wavelength range $[600, 1000]$ nm. The concentration of the various species in the background have been summarized in Table 6.1. The extinction coefficients for the species have been found in the literature [35]. We take the value [17] of $\Psi = 9.4$ based on the wavelength of 600 nm and the prefactor $b = 1.4$ in Equation (2.9).

Species	HbO_2	HbR	H_2O	Lipids
Units	μM	μM	%	%
Background	17	7	0.15	0.6
Tumor	25	15	0.25	0.5

Table 6.1: Concentration of different species in the background and the tumor.

6.1. Forward solver. For the preconditioner we choose an incomplete LU factorization implemented using SuperLU [28] and we considered the parameters `fill_factor` $\in \{5, 10, 15\}$ and `drop_tol` $\in \{10^{-3}, 10^{-4}, 10^{-5}\}$. We consider the following transformation of the linear systems (3.1)

$$\left(\mathbf{K} + \bar{\sigma}'\mathbf{R} + \sigma_j\mathbf{M} + (\sigma_j' - \bar{\sigma}')\mathbf{R}\right)\mathbf{x}_j = \mathbf{b}$$

Varying problem size, drop tol $= 10^{-4}$, fill factor $= 10$									
Dim k	$N = 16,271$			$N = 52,425$			$N = 87,431$		
	Iter	MVP [s]	Tot. [s]	Iter	MVP [s]	Tot. [s]	Iter	MVP [s]	Tot. [s]
0	2927	27.90	41.84	3992	146.39	238.11	5426	388.35	740.34
5	2164	18.09	32.84	3028	124.39	232.86	4922	315.79	495.68
10	2093	18.94	34.29	2845	116.44	195.03	4768	300.87	469.59
15	2072	19.34	35.21	2850	103.59	180.86	4709	311.00	479.408
Varying fill factor, $N = 52,425$, drop tol; $= 10^{-4}$									
	fill factor $= 5$			fill factor $= 10$			fill factor $= 15$		
0	5080	129.91	297.57	3992	146.39	238.11	3688	195.37	262.99
5	3710	105.21	222.64	3028	124.39	232.86	2809	165.09	248.91
10	3614	98.20	210.50	2845	116.44	195.03	2664	145.67	218.17
15	3729	91.93	197.10	2850	103.59	180.86	2651	131.70	194.64
Varying drop tol, fill factor $= 10$, $N = 52,425$									
	drop tol $= 10^{-3}$			drop tol $= 10^{-4}$			drop tol $= 10^{-5}$		
0	3485	99.61	179.86	3992	146.39	238.11	5293	193.75	314.37
5	2665	77.21	169.98	3028	124.39	232.86	3859	143.45	248.03
10	2541	75.81	161.72	2845	116.44	195.03	3721	155.99	259.73
15	2505	71.76	148.10	2850	103.59	180.86	3744	144.06	228.62

Table 6.2: Summary of augmented GMRES solver for different problem sizes, and preconditioner parameters. The number k refers to the dimension of the deflation space. All systems were solved till it converged to a relative tolerance of 10^{-8}. As can be seen, on average $10-20\%$ improvement was observed for all systems in terms of total computation time, by using deflation.

for $j = 1, \ldots, N_\lambda$ and $\bar{\sigma}'$ is the mean of σ'_j. We then define $\mathbf{K} \leftarrow \mathbf{K} + \bar{\sigma}'\mathbf{R}$ and $\sigma'_j \leftarrow \sigma'_j - \bar{\sigma}'$. This transformation essentially the leaves the solution unchanged but improves the convergence of our solver, since the modified matrix $\mathbf{K} + \bar{\sigma}'\mathbf{R}$ contains average information about the refractive index mismatch boundary conditions. Other transformations involving the minimum or maximum over σ'_j may also be considered.

We now report the results of our solver on a variety of test problem sizes, and preconditioner parameter such as fill factor and drop tolerance. The column labeled 'Iter' reports the total number of iterations across 100 wavelengths, 'MVP [s]' reports the CPU time spent on matrix-vector products as well as application of the preconditioner, and finally 'Tot. [s]' reports the total CPU run time of Algorithm 1 (including the pre-computation time for computing the initial guess $\tilde{\mathbf{x}}_{0,j}$ and the matrices $\mathbf{U}, \mathbf{C}$ obtained by solving the shift-invariant system and generating the augmented space for each system $\mathbf{U}_j$ and $\mathbf{C}_j$). From the table it can be seen that by increasing the dimension of the deflation space k, the total number of iterations decrease but the cost per iteration increases as a result of extra orthogonalization w.r.t. $\mathbf{C}_j$. Therefore, there is a trade-off between the number of iterations and total run time and adding additional vectors in the deflation space is a case of diminishing returns. For the range of parameters we experimented with, typically $k = 5, 10$ produces the best results in terms of total CPU time. However, for problem sizes larger than we are considering, the cost of matrix-vector products may be the dominant cost so that it might be beneficial to use a larger deflation space.

We would like to emphasize that care should be taken to interpret the results in Table 6.2. The algorithm has been implemented in Python (which is an interpreted language) and because it uses pre-compiled code for parts of the computation, the timing results may be slightly different if the entire algorithm were implemented in a single programming language. In particular, we expect the overall computation time would be lower if using a compiled language such as C/C++ and the performance gains from our algorithm to be higher.

6.2. Compression. We now discuss the results of the compression scheme presented in Section 4. We consider the same geometry that was used in the previous subsection. The sources are placed on the top of the phantom whereas the receivers are placed on the bottom. For each source, there are 9

detectors constrained to move along with the source that are evenly placed co-axially with the source with 0.5 cm distance from each other. The number of sources varied from $N_s = 4, \ldots, 25$ and the number of wavelengths vary between $N_\lambda = 11, \ldots, 81$. With these parameters, the maximum number of parameters are $18,225$. The discretized grid has $N = 52,425$ degrees of freedom.

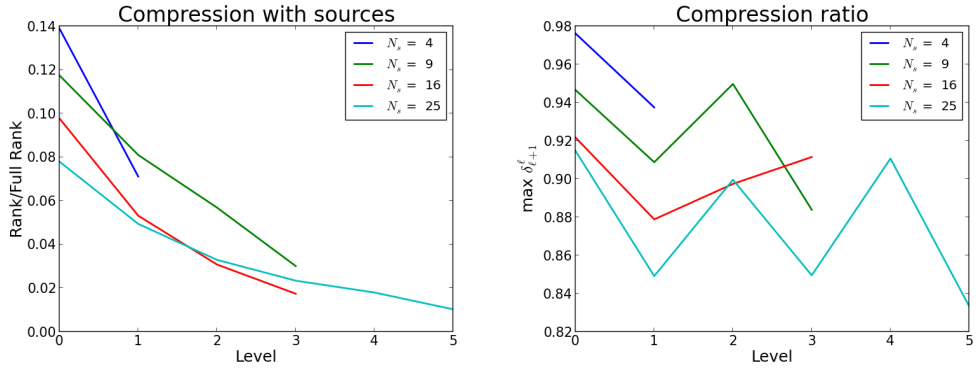

Fig. 6.1: (left) Compression defined as Rank/Full rank of the recursive SVD computed as a function of the level of the tree, with increasing number of sources. Here Rank is defined as the number of columns of $\mathbf{U}$ and Full Rank $= \min\{M, N\}$. (right) The compression ratio $\delta^\ell_{\ell+1} \stackrel{\text{def}}{=} r_\ell/(2r_{\ell+1})$ as function of the level. As can be seen, on average, there is higher levels of compression at higher levels in the tree. Here, level 0 is the root of the tree. A tolerance of 10^{-6} was used for the truncation. Furthermore, $N_{ds} = 9$ and $N_\lambda = 41$.

In the examples, we will describe, we have used the randomized SVD for both computing the low-rank factorization at the leaf level and to compute the agglomeration of the low-rank factors as we go up the tree. Similar results are obtained using the partially pivoted Adaptive Cross Approximation and will not be displayed here. We first consider the compression by varying the number of sources and keeping all other parameters fixed. We assume that $N_{ds} = 9$ and $N_\lambda = 41$ and a tolerance of 10^{-6} was used for truncating the rank of the sub-blocks and the agglomeration. As can be seen, at higher levels in the tree (closer to the root) we observe a higher level of compression because there is a greater redundancy of information globally as opposed to locally. Furthermore, with increasing number of sources we observe a higher level of compression which implies that there is redundancy both in terms of wavelengths and the source-detector positions. The results are displayed in Figure 6.1. We also plot the maximum compression ratio $\delta^\ell_{\ell+1} \stackrel{\text{def}}{=} r_\ell/(2r_{\ell+1})$ (computed across all the nodes at level ℓ) which is the ratio of the ranks at level ℓ obtained by agglomerating 2 sub-blocks at level $\ell + 1$ with ranks $\ell + 1$. As can be seen, $\delta^\ell_{\ell+1} < 1$ at all levels indicating that there is compression, not only at the leaves, but compression at every level in the tree. This justifies using a hierarchical compression scheme and the cost analysis performed in Section 4.5.

Next we compute the compression as a function of tree level with varying tolerance used to truncate the ranks of the sub-blocks. The results are presented in Figure 6.2. As can be seen with a higher tolerance the ranks at each level decreases dramatically. Here, we have fixed $N_{ds} = 9$, $N_\lambda = 41$ and $N_s = 9$. We also present results of computing the compression as a function of tree level with varying number of wavelengths used to illuminate the object. All other parameters are fixed as $N_{ds} = 9$, tol $= 10^{-6}$ and $N_s = 9$. The results are also presented in Figure 6.2. We can see that with increasing number of wavelengths there is a higher level of compression at higher levels in the tree.

Finally, we compare the run time of the recursive SVD algorithm proposed in Section 4 with RandSVD applied to the entire measurement operator. As can be seen, the computational time for the recursive algorithm is far lower and is therefore, more efficient.

6.3. Reconstruction results. In the following experiments we will use the following metrics to measure the error in the shape perturbation. In Experiment 1 and 2 we use the full Born operator to generate the measurements, however, we use the compressed operator in the reconstructions. In

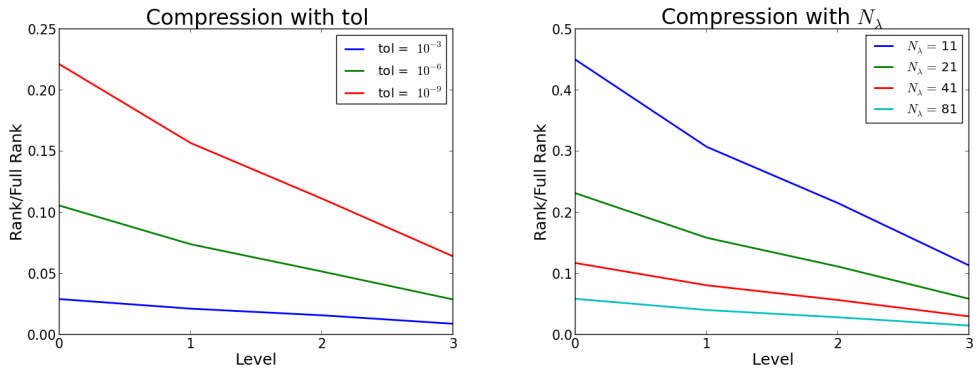

Fig. 6.2: (left) Compression defined as Rank/Full rank of the recursive SVD computed as a function of the level of the tree, with varying tolerance used to compress the low-rank factors. Here $N_{ds} = 9$, $N_\lambda = 41$ and $N_s = 9$. (right) The compression of the recursive SVD computed as a function of the level of the tree, with varying N_λ. Here $N_{ds} = 9$, tol $= 10^{-6}$ and $N_s = 9$.

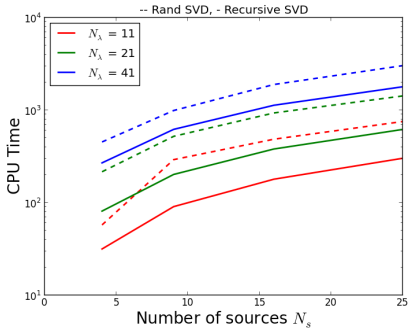

Fig. 6.3: Comparison of time taken to factorize the Born operator $\mathbf{H}$ to a tolerance of 10^{-6} with N_s ranging from 4 to 25 and N_{ds} fixed at 9. We compare the CPU timing of RandSVD in Algorithm 5 applied to the entire matrix directly with the Recursive SVD algorithm proposed in Section 4. The time for constructing $\mathbf{H}$ is not included. The grid size was $16,721$.

Experiment 3 we use measurements from the full diffusion equation and use the compressed Born model in the reconstruction. As a result, we avoid Let $\boldsymbol{\mu}$ denote the true shape perturbation and let $\hat{\boldsymbol{\mu}}$ denote the recovered shape perturbation. The first is standard relative L_2 error $\|\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}\|_2 / \|\boldsymbol{\mu}\|_2$. For piecewise constant medium, a different metric known as the Dice coefficient (see for example, [27]) provides more information about localization of the perturbation. It can be defined as

$$D(\boldsymbol{\mu}, \hat{\boldsymbol{\mu}}) \stackrel{\text{def}}{=} 2 \frac{|\boldsymbol{\mu} \bigcap \hat{\boldsymbol{\mu}}|}{|\boldsymbol{\mu}| + |\hat{\boldsymbol{\mu}}|}$$

where $|\boldsymbol{\mu} \bigcap \hat{\boldsymbol{\mu}}|$ corresponds to the number of non-zero pixels that are contained in both the true shape $\boldsymbol{\mu}$ and the reconstructed shape $\hat{\boldsymbol{\mu}}$ whereas $|\boldsymbol{\mu}|$ and $|\hat{\boldsymbol{\mu}}|$ correspond to the number of non-zero pixels in the true and the reconstructed shapes respectively. In order to If the object is recovered fully, then the Dice coefficient would be 1.

Experiment 1: In this experiment we study how the accuracy of the low-rank representation for $\mathbf{H}$ affects the error in the reconstruction. We consider $N_s = 4$ sources and $N_{ds} = 9$ detectors/source and $N_\lambda = 25$ wavelengths totaling 900 measurements generated using the Born model but with the full measurement operator $\mathbf{H}$. We only consider reconstruction of the shape perturbation and consider

the concentration of the chromophores as known and provided in Table 6.1. Gaussian noise is added such that the signal-to-noise (SNR) ratio is 33 dB to simulate observational noise. The SNR is defined as

$$\text{SNR} = 20 \log_{10} \frac{\|\mathbf{y}\|_2}{\|\boldsymbol{\eta}\|_2}$$

where the noise $\boldsymbol{\eta}$ has been defined in Equation (2.12). The grid size is chosen to be $52,425$. A "true" shape perturbation is obtained by using three randomly generated basis functions which leads to $N_p = 15$. We report the rank of the measurement operator as a function of the global tolerance used for low-rank representation and the error in the reconstruction using the low-rank operator $\hat{\mathbf{H}}$. The results are reported in Table 6.3. As can be seen that for very small tolerance, the error in the reconstruction is hardly noticeable. Therefore, a larger rank of the low-rank representation $\hat{\mathbf{H}}$ does not affect the reconstruction error below a certain tolerance and therefore the compressed operator $\hat{\mathbf{H}}$ can be used as a surrogate for the full matrix $\mathbf{H}$ with little or no loss in accuracy in the reconstruction. When the number of measurements is large, this can represent significant savings in computational time.

tol	10^{-3}	10^{-6}	10^{-9}
Rank	48	120	288
L_2err.	40.62%	40.52%	40.52%
Dice	0.82	0.827	0.827
Time [s]	2.20	2.23	2.32

Table 6.3: Rank of the compressed operator and the error in the reconstruction of the shape perturbation as a function of tolerance used to compress $\mathbf{H}$. See experiment 1 for more details. 'Time [s]' indicates the CPU time of solving the optimization problem with the compressed operator. See Experiment 1 for more details.

Experiment 2: In this experiment we study the reconstruction of chromophore concentrations as well as the shape perturbation. The number of measurements and the process of generating them is the same as Experiment 1 except with SNR 30 dB. For the reconstruction, the radial basis functions were randomly initialized and a truncated measurement operator $\hat{\mathbf{H}}$ computed using tolerance 10^{-6} was used in the reconstruction. The error of the reconstruction of the chromophore concentrations as well the relative L_2 error of the shape perturbation are reported in Table 6.4. As can be seen from the Table, the added difficulty in recovering the shape perturbation as well as the chromophore concentrations affects the reconstruction error of the shape perturbation slightly. Moreover, the concentrations of the chromophore species are recovered fairly accurately. This is consistent with the observations in [25, 26, 27].

Species	HbO_2	HbR	H_2O	Lipids	L_2 err.	Dice
Recon.	1.6%	0.2%	4.9%	4.4%	52.72%	0.79

Table 6.4: Error in the reconstruction of different species and the shape perturbation corresponding to Experiment 2.

Experiment 3: In this example we examine the validity of the Born approximation. As mentioned earlier, experimental studies have validated the utility of the Born approximation (see for e.g. [27]). However, in order to prevent committing an "inverse crime" we use data generated from the full diffusion equation, Equation (2.1) solving using finite elements with the same boundary conditions and use it to reconstruct both the chromophore concentrations and the shape parameter. The comparison between the measurements generated using the Born model and those generated from the full diffusion equation is provided in Figure 6.5. In addition a noise of 50 dB was added to simulate

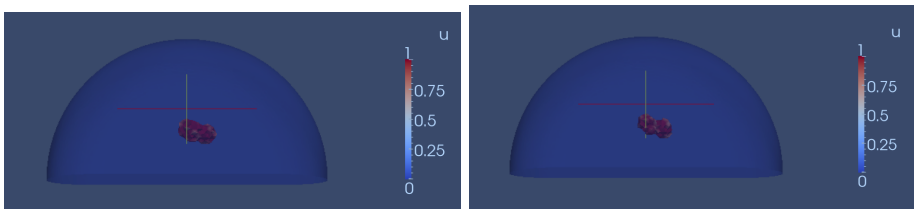

Fig. 6.4: Reconstruction in Experiment 2 (left) true anomaly and (right) reconstruction. The error is described in Table 6.4.

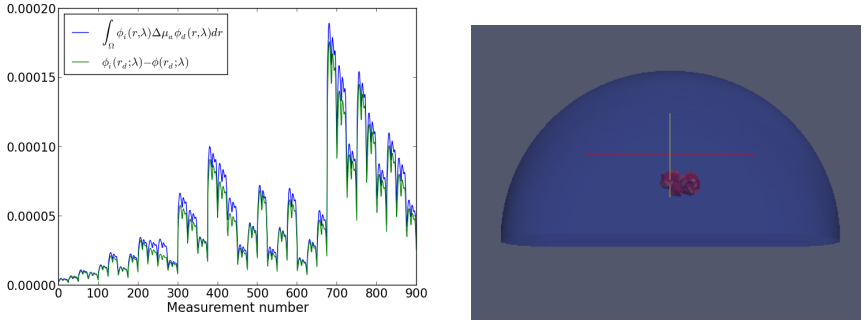

Fig. 6.5: (left) Comparison of the measurements generated using the Born model and that generated from the full diffusion equation. As can be seen, the agreement between the measurements is quite close. (right) Reconstruction of the shape perturbation using measurements from full diffusion equation. The errors are reported in Table 6.5.

observational error in realistic conditions. All other settings were the same as that in Experiment 2. The error of the reconstruction of the chromophore concentrations as well the L_2 error of the shape perturbation are reported in Table 6.5. The error in the reconstruction is higher than that obtained from Experiment 2. Since the full diffusion equation was used in generating the measurements, there is a modeling error which corresponds to about 20 dB SNR and because of this, the Born model cannot exactly recover the shape perturbation. However, the reconstruction is still satisfactory as can be seen from the Figure 6.5.

Species	HbO_2	HbR	H_2O	Lipids	L_2 err.	Dice
Recon.	5.1%	1.2%	1.2%	9.3%	71.80%	0.65

Table 6.5: Error in the reconstruction of different species and the shape perturbation corresponding to Experiment 3. Measurements were generated from the full diffusion equation, i.e., Equation (2.1). See also Figure 6.5.

7. Conclusions and future work. We have presented a fast algorithm for recovering shape of the perturbations and chromophore concentrations that is composed of three parts - a fast Krylov subspace approach for accelerating the solution of the incident and the adjoint field for multiple wavelengths, constructing a low-rank approximation to the sensitivity matrix $\mathbf{H}$ using an approach that accounts for redundancies across wavelengths which is then recursively combined across source-detectors pairs, and finally solving an optimization problem for recovering parameters with the low-rank approximation to $\mathbf{H}$. The results indicate that significant gains can be obtained both in terms of computational costs and storage requirements. We provide detailed numerical experiments that validates our claims and further provide a detailed analysis of the computational costs and the error.

The algorithms were demonstrated on a challenging synthetic inversion case on a complex geometry which highlights the flexibility of our algorithms.

Future work includes extension of our algorithms to the fully nonlinear setting. In this setting, computation of the incident ϕ_i and adjoint fields ϕ_d and the construction of the measurement operator $\mathbf{H}$ for the Born approximation, must be performed at every Newton or Gauss-Newton step. We therefore, believe that adopting the ideas proposed in this paper will be even more beneficial in the fully nonlinear case. Other possible extensions include a recycling strategy of the kind described in Section 3 to multiple wavelengths and multiple right hand sides and to the full diffusion equation, Equation (2.1). Additional work is currently underway in our lab to compare the reconstructions obtained the synthetic case with real data obtained from experiments. Based on previous work [25, 26, 27], we anticipate that the reconstructions will indeed be excellent even in the hyperspectral case.

8. Acknowledgements. We would like to thank Nishanth Krishnamurthy, Pami G. Anderson, Jana Kainerstorfer, and Angelo Sassaroli for useful discussions. The first author would also like to thank Tania Bakhos with her help in generating the meshes. This work was supported by NIH Grant R01-CA154774. Additionally, the second author was supported by NSF Grant DMS 1217161.

Appendix A. Computing low-rank representations.

A.1. Randomized SVD. Randomized algorithms for matrix decompositions were derived in a series of papers [29, 32, 19]. In this work, we consider the algorithms described in [19], which was also used by [9]. Suppose we wish to compute the rank r decomposition of the matrix. The algorithm begins by computing a matrix $\mathbf{Q}$ that approximates the column space of $\mathbf{A}$ as

$$\|\mathbf{Q}\mathbf{Q}^*\mathbf{A} - \mathbf{A}\| \leq \epsilon \tag{A.1}$$

where ϵ is a user-defined tolerance. The matrix $\mathbf{Q}$ is obtained by computing the matrix-vector products of $\mathbf{A}$ with the matrix $\mathbf{\Omega}_1 \in \mathbb{R}^{n\times(r+p)}$, with i.i.d. entries drawn from a standard normal distribution, $\mathcal{N}(0,1)$, and then computing a basis for the resulting matrix using QR or the SVD. Here, p is an oversampling parameter that is chosen a priori and typically $p \sim 20$. A matrix $\mathbf{Q}$ that satisfies the bound (A.1) can be converted into a low-rank representation using matrix manipulations. A discussion of the choice of oversampling factor and the low-rank conversion is described in [19].

The algorithm has good performance when the singular values of the matrix $\mathbf{A}$ decay rapidly. If this is not the case, the power method is applied to improve the convergence of the algorithm. Since the rank of the matrix is not known a priori, we use an adaptive approach to estimate the range that is combined with an error estimator described in [19]. The algorithm is summarized in 5. The dominant computational cost is computing the matrix-vector products with the matrix $\mathbf{A}$ which costs $\mathcal{O}(mn(r+p))$. The total cost of computing the low-rank representation is $\mathcal{O}\left(mn(r+p) + r^2(m+n)\right)$.

A.2. Adaptive Cross Approximation. The idea behind the cross approximation is based on the result described in [4], which states that supposing a matrix A is well approximated by a low-rank matrix, by a clever choice of k columns indexed as $\mathcal{J}$ and k rows indexed as $\mathcal{I}$ of the matrix $\mathbf{A}$, we can approximate $\hat{\mathbf{A}}$ of the form

$$\|\mathbf{A} - \hat{\mathbf{A}}\| \leq \varepsilon \qquad \hat{\mathbf{A}} = \mathbf{A}(:,\mathcal{J})\mathbf{A}(\mathcal{I},\mathcal{J})^{-1}\mathbf{A}(\mathcal{I},:)$$

This decomposition relies on a result from [14], which states that if there is a sufficiently good low rank approximation to a matrix, then there exists a cross-approximation with almost the same approximation quality.

Algorithm 6 describes a simple heuristic to compute such a cross approximation that is based on successive approximations by rank-1 matrices. It has the property that if the matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$ has an exact rank $r < \min\{m,n\}$, this algorithm will terminate in r steps and defining

$$\mathbf{S}_r = \sum_{k=1}^{r} \mathbf{u}_k \mathbf{v}_k^T$$

Algorithm 5 Randomized SVD [19]

$r = 1$ (initial guess for rank), $k = 10$ (error estimator), $p = 20$ (oversampling factor)
while true do
  Contruct matrices $\mathbf{\Omega}_1 \in \mathbb{R}^{n\times(r+p)}$ and $\mathbf{\Omega}_2 \in \mathbb{R}^{n\times k}$ with entries drawn from i.i.d. normal distribution $\mathcal{N}(0,1)$.
  $\mathbf{Y} = \mathbf{A}\mathbf{\Omega}_1$
  $[\mathbf{U}_r, \mathbf{S}_r, \mathbf{V}_r] = \text{SVD}(\mathbf{Y})$
  $\mathbf{Q} = \mathbf{U}(:, 1:r)$
  $\mathbf{B} \overset{\text{def}}{=} \mathbf{Q}^T\mathbf{A}$
  Error estimate $e_r = \|\mathbf{A}\mathbf{\Omega}_2 - \mathbf{Q}\mathbf{Q}^T\mathbf{A}\mathbf{\Omega}_2\|$
  if $e_r > \varepsilon\mathbf{S}_r(1,1)$ **then**
    Increment r
  else
    $[\mathbf{U}_1, \mathbf{S}, \mathbf{V}]$ =SVD($\mathbf{B}$)
    Break
  end if
end while
return $[\mathbf{U}, \mathbf{S}, \mathbf{V}]$ where $\mathbf{U} = \mathbf{U}_r\mathbf{U}_1$.

Algorithm 6 Cross Approximation using Full Pivoting [5]

Initialize

$$\mathbf{R}_0 = \mathbf{A}, \qquad \mathbf{S} = \mathbf{0}$$

for all $k = 0, 1, 2, \ldots$ **do**
  $(i^*_{k+1}, j^*_{k+1}) := \arg\max_{i,j}|(\mathbf{R}_k)_{ij}|$ and $\gamma_{k+1} = \left(\mathbf{A}_{i^*_{k+1}, j^*_{k+1}}\right)^{-1}$.
  if $\gamma_{k+1} \neq 0$ **then**
    Compute column $\mathbf{u}_{k+1} := \gamma_{k+1}\mathbf{R}_k\mathbf{e}_{j_{k+1}}$ and row $\mathbf{v}_{k+1} := \mathbf{R}_k^T\mathbf{e}_{i_{k+1}}$
    New residue and approximation

$$\mathbf{R}_{k+1} := \mathbf{R}_k - \mathbf{u}_{k+1}\mathbf{v}_{k+1}^T \qquad \mathbf{S}_{k+1} := \mathbf{S}_k + \mathbf{u}_{k+1}\mathbf{v}_{k+1}^T$$

  else
    Terminate algorithm with exact rank $k-1$
  end if
end for

we have that $\mathbf{S}_r = \mathbf{A}$ in exact arithmetic. Furthermore, it exactly reproduces the r pivot rows and columns of $\mathbf{A}$. Of course, the principal disadvantage of this algorithm is that, to generate a rank-k approximation, it requires $\mathcal{O}(rmn)$ steps, which is not feasible for large matrices. The bottleneck arises from calculating the pivot indices (i^*_k, j^*_k) which requires generating all the entries of the matrix $\mathbf{A}$.

Several heuristic strategies have been proposed to reduce the complexity of the fully pivoting cross approximation algorithm. In particular, one such algorithm is called partially pivoted Adaptive Cross Approximation algorithm that has a complexity $\mathcal{O}(r^2(m+n))$. A practical version of the algorithm, which includes a termination criteria based on an heuristic approximation to the relative approximation in the Frobenius norm, can be found in [5]. This is the version we will use in the rest of the paper.

REFERENCES

[1] Alireza Aghasi, Misha Kilmer, and Eric L Miller. Parametric level set methods for inverse problems. *SIAM Journal on Imaging Sciences*, 4(2):618–650, 2011.

[2] Simon Arridge and John Schotland. Optical tomography: forward and inverse problems. *arXiv preprint arXiv:0907.2586*, 2009.

[3] Simon R Arridge. Optical tomography in medical imaging. *Inverse problems*, 15(2):R41, 1999.

[4] Mario Bebendorf. Approximation of boundary element matrices. *Numerische Mathematik*, 86(4):565–589, 2000.

[5] Mario Bebendorf and Sergej Rjasanow. Adaptive low-rank approximation of collocation matrices. *Computing*, 70(1):1–24, 2003.

[6] Peter Benner, Serkan Gugercin, and Karen Willcox. A survey of model reduction methods for parametric systems. Preprint MPIMD/13-14, Max Planck Institute Magdeburg, 2013.

[7] David Boas. A fundamental limitation of linearized algorithms for diffuse optical tomography. *Optics Express*, 1(13):404–413, 1997.

[8] Gregory Boverman, Eric L Miller, Ang Li, Quan Zhang, Tina Chaves, Dana H Brooks, and David A Boas. Quantitative spectroscopic diffuse optical tomography of the breast guided by imperfect a priori structural information. *Physics in medicine and biology*, 50(17):3941, 2005.

[9] Stéphanie Chaillat and George Biros. FaIMS: A fast algorithm for the inverse medium problem with multiple frequencies and multiple sources for the scalar Helmholtz equation. *Journal of Computational Physics*, 231(12):4403–4421, 2012.

[10] Tony F Chan and Michael K Ng. Galerkin projection methods for solving multiple linear systems. *SIAM Journal on Scientific Computing*, 21(3):836–850, 1999.

[11] Paul G Constantine, David F Gleich, and Gianluca Iaccarino. Spectral methods for parameterized matrix equations. *SIAM Journal on Matrix Analysis and Applications*, 31(5):2681–2699, 2010.

[12] Eric de Sturler and Misha E Kilmer. A regularized Gauss-Newton trust region approach to imaging in diffuse optical tomography. *SIAM Journal on Scientific Computing*, 33(5):3057–3086, 2011.

[13] Christophe Geuzaine and Jean-François Remacle. Gmsh: A 3-d finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11):1309–1331, 2009.

[14] Sergei A Goreinov, Eugene E Tyrtyshnikov, and Nickolai L Zamarashkin. A theory of pseudoskeleton approximations. *Linear Algebra and Its Applications*, 261(1):1–21, 1997.

[15] Dirk Grosenick, K Thomas Moesta, Michael Möller, Jörg Mucke, Heidrun Wabnitz, Bernd Gebauer, Christian Stroszczynski, Bernhard Wassermann, Peter M Schlag, and Herbert Rinneberg. Time-domain scanning optical mammography: I. recording and assessment of mammograms of 154 patients. *Physics in medicine and biology*, 50(11):2429, 2005.

[16] Dirk Grosenick, Heidrun Wabnitz, K Thomas Moesta, Jörg Mucke, Michael Möller, Christian Stroszczynski, Jana Stößel, Bernhard Wassermann, Peter M Schlag, and Herbert Rinneberg. Concentration and oxygen saturation of haemoglobin of 50 breast tumours determined by time-domain optical mammography. *Physics in medicine and biology*, 49(7):1165, 2004.

[17] Dirk Grosenick, Heidrun Wabnitz, K Thomas Moesta, Jörg Mucke, Peter M Schlag, and Herbert Rinneberg. Time-domain scanning optical mammography: Ii. optical properties and tissue parameters of 87 carcinomas. *Physics in medicine and biology*, 50(11):2451, 2005.

[18] Wolfgang Hackbusch. New estimates for the recursive low-rank truncation. *Preprint number: 34, Max-Planck-Institut f'ur Mathematik in den Naturwissenschaften Leipzig*, 2014.

[19] N. Halko, P.G. Martinsson, and J.A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.

[20] Thomas JR Hughes. *The finite element method: linear static and dynamic finite element analysis*. Courier Dover Publications, 2012.

[21] Damon Hyde, Misha Kilmer, Dana H Brooks, and Eric Miller. Analysis and exploitation of matrix structure arising in linearized optical tomographic imaging. *SIAM Journal on Matrix Analysis and Applications*, 29(4):1065–1082, 2007.

[22] Ilse CF Ipsen, CT Kelley, and SR Pope. Rank-deficient nonlinear least squares problems and subset selection. *SIAM Journal on Numerical Analysis*, 49(3):1244–1266, 2011.

[23] Misha E Kilmer and Eric de Sturler. Recycling subspace information for diffuse optical tomography. *SIAM Journal on Scientific Computing*, 27(6):2140–2166, 2006.

[24] Daniel Kressner and Christine Tobler. Low-rank tensor Krylov subspace methods for parametrized linear systems. *SIAM Journal on Matrix Analysis and Applications*, 32(4):1288–1316, 2011.

[25] Fridrik Larusson, Pamela G Anderson, Elizabeth Rosenberg, Misha E Kilmer, Angelo Sassaroli, Sergio Fantini, and Eric L Miller. Parametric estimation of 3D tubular structures for diffuse optical tomography. *Biomedical optics express*, 4(2):271, 2013.

[26] Fridrik Larusson, Sergio Fantini, and Eric L Miller. Hyperspectral image reconstruction for diffuse optical tomography. *Biomedical optics express*, 2(4):946, 2011.

[27] Fridrik Larusson, Sergio Fantini, and Eric L Miller. Parametric level set reconstruction methods for hyperspectral diffuse optical tomography. *Biomedical optics express*, 3(5):1006–1024, 2012.

[28] Xiaoye S. Li and Meiyue Shao. A supernodal approach to imcomplete LU factorization with partial pivoting. *ACM Trans. Mathematical Software*, 37(4), 2010.

[29] E. Liberty, F. Woolfe, P.G. Martinsson, V. Rokhlin, and M. Tygert. Randomized algorithms for the low-rank approximation of matrices. *Proceedings of the National Academy of Sciences*, 104(51):20167–20172, 2007.

[30] Anders Logg and Garth N Wells. Dolfin: Automated finite element computing. *ACM Transactions on Mathematical Software (TOMS)*, 37(2):20, 2010.

[31] Vadim A Markel, Vivek Mital, and John C Schotland. Inverse problem in optical diffusion tomography. iii. inversion formulas and singular-value decomposition. *JOSA A*, 20(5):890–902, 2003.

[32] Per-Gunnar Martinsson, Vladimir Rokhlin, and Mark Tygert. A randomized algorithm for the decomposition of matrices. *Applied and Computational Harmonic Analysis*, 30(1):47–68, 2011.

[33] Luís Augusto Motta Mello, Eric De Sturler, Glaucio H Paulino, and Emílio Carlos Nelli Silva. Recycling krylov subspaces for efficient large-scale electrical impedance tomography. *Computer Methods in Applied Mechanics and Engineering*, 199(49):3101–3110, 2010.

[34] Michael L Parks, Eric De Sturler, Greg Mackey, Duane D Johnson, and Spandan Maiti. Recycling Krylov subspaces for sequences of linear systems. *SIAM Journal on Scientific Computing*, 28(5):1651–1674, 2006.

[35] Scott Prahl. Tabulated molar extinction coefficient for hemoglobin in water. *http://omlc.org/spectra/hemoglobin/summary.html*.

[36] M Schweiger and SR Arridge. Optical tomographic reconstruction in a complex head model using a priori region boundary information. *Physics in Medicine and Biology*, 44(11):2703, 1999.

[37] Natasha Shah, David Hsiang, Bruce J Tromberg, John Butler, Albert E Cerussi, and Dorota Jakubowski. Spatial variations in optical and physiological properties of healthy breast tissue. *Journal of biomedical optics*, 9(3):534–540, 2004.

[38] V. Simoncini and D.B. Szyld. Recent computational developments in Krylov subspace methods for linear systems. *Numerical Linear Algebra with Applications*, 14(1):1–59, 2007.

[39] Shun Wang, Eric de Sturler, and Glaucio H Paulino. Large-scale topology optimization using preconditioned krylov subspace methods with recycling. *International Journal for Numerical Methods in Engineering*, 69(12):2441–2468, 2007.

[40] Yao Yue and Karl Meerbergen. Accelerating optimization of parametric linear systems by model order reduction. *SIAM Journal on Optimization*, 23(2):1344–1370, 2013.